R. NOBLE

Proceedings from the

1st $\mu^+ \mu^-$ Collider

Workshop

Napa, California

1992

Published in Nuclear Instruments and Methods

1994

The UCLA Center for Advanced Accelerators

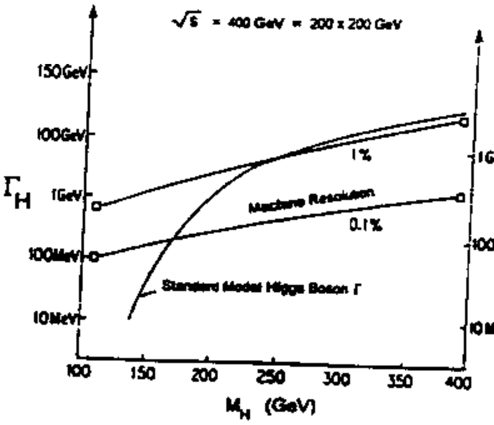

Fig. 4. Higgs search at a $\mu^+\mu^-$ collider. Required machine resolution and the expected Higgs width.

using a broad energy sweep. The corresponding cross section is small (see Figs. 3 and 5). Once an approximate mass is determined, a strategy for the energy sweep through the resonance can be devised. The study of the t quark through $t\bar{t}$ production would also be interesting.

Finally, another possibility is to use the polarization of the $\mu^+\mu^-$ particles orientated so that only scalar interactions are possible (thus eliminating the background from single photon intermediate states as shown in Fig. 3) [6]. However, there would be a trade-off with luminosity and thus a strategy would have to be devised to maximize the possibility of success in the energy sweep through the resonance.

At the Napa workshop the possibility of developing a $\mu^+\mu^-$ collider in the 10^{31} cm^{-2} s^{-1} region was considered and appears feasible. It is less certain that the high energy resolution required for the Higgs sweep can be obtained. We summarize in Table 1 some of the key issues in the Higgs search.

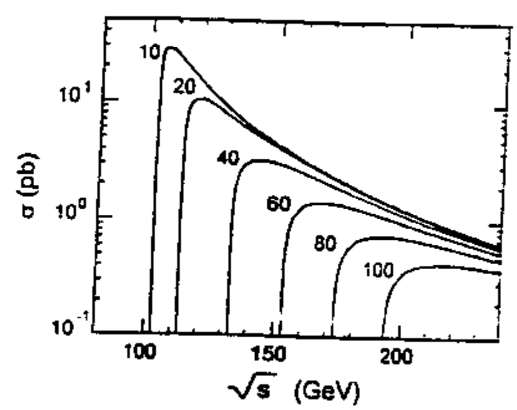

Fig. 5. Total cross section for $\mu^+\mu^- \to \phi^0 Z$ as a function of $\sqrt{s}$ for the fixed $m\phi_0$ values indicated by the numbers (in GeV) beside each line. (Adapted from Ref. [4].)

Table 1
Key issues in the Higgs search

There is growing evidence that one Higgs particle is below 2 M_Z
SUSY Higgs – 3 Higgs – one near M_Z (possibly up to ~130 GeV); extremely hard to detect
Hadron machines can search for these Higgs provided: i $\int \mathcal{L}\, dt \geq 10^5$ pb (LHC); ii the background for $H \to \gamma\gamma$ is small enough
A $\mu^+\mu^-$ collider with $\mathcal{L} \gtrsim 10^{31}$ cm^{-2} s^{-1} operating between 100–180 GeV could discover the Higgs in 2000+ provided sufficient energy resolution is achieved

$\mu^+\mu^-$ colliders could also be very important in the TeV energy range; however, since the cross sections for new particle production are much smaller, the luminosity requirements would be $\mathcal{L}_{\mu^+\mu^-} \geq 10^{33}$ cm^{-2} s^{-1}. This is the energy range where e^+e^- linear colliders are extremely difficult to develop [1]. This possibility will be the subject of a second Napa workshop to be held in 1994.

Acknowledgement

I wish to thank the participants of the workshop, and in particular Dave Neuffer and Andy Sessler, Bill Barletta and Kirk McDonald for discussions.

References

[1] M. Tigner, presented at the 3rd Int. Workshop on Advanced Accelerator Concepts, Port Jefferson, NY, (1992); other paper to be published in AIP, 1992, ed. J. Wurtle.

[2] Early references for $\mu\mu$ colliders are: E.A. Perevedentsev and A.N. Skrinsky, Proc. 12th Int. Conf. on High Energy Accelerators eds. F.T. Cole and R. Donaldson, 1983, p. 485; D. Neuffer, Proc. 12th Int. Conf. on High Energy Accelerators, eds F.T. Cole and R. Donaldson, 1983, p. 481; D. Neuffer, Particle Accel. 14 (1983) 75; D. Neuffer, in: Advanced Accelerator Concepts, AIP Conf. Proc. 156 (1987) 201.

[3] Reports from the Napa Workshop: D. Neuffer, $\mu^+\mu^-$ Colliders: Possibilities and Challenges; W. Barletta and A. Sessler, Characteristics of a high energy $\mu^+\mu^-$ collider based on electro-production of muons.

[4] See for primary references to the theoretical estimates here S. Dawson, J.F. Gunion, H.E. Haber and G.L. Kane, The Physics of the Higgs Bosons: Higgs Hunter's Guide (Addison Wesley, Menlo Park, 1989).

[5] For example the Compact Muon Solenoid detector at the LHC, CERN reports.

[6] This possibility came up in discussion with K. McDonald, Princeton.

Nuclear Instruments and Methods in Physics Research A 350 (1994) 27–35
North-Holland

INSTRUMENTS
& METHODS
IN PHYSICS
RESEARCH
Section A

μ^+–μ^- colliders: possibilities and challenges ☆

David V. Neuffer

Continuous Electron Beam Accelerator Facility, 12000 Jefferson Avenue, Newport News, VA 23606, USA

Received 8 April 1994

The current status of the μ^+–μ^- collider concept is reviewed and discussed. In a reference scenario, a high-intensity pulsed proton accelerator (of K-factory class) produces large numbers of secondary π's in a nuclear target, which produce muons by decay. The muons are collected and cooled (by "ionization cooling") to form high-intensity bunches that are accelerated to high-energy collisions. High-luminosity μ^+–μ^- and μ^-–p colliders at TeV or higher energy scales may be possible. Challenges in implementing the scenario are described. Possible variations in muon production, accumulation, and collisions are discussed; further innovations and improvements are encouraged.

1. Introduction

Current and planned highest-energy colliders are hadronic proton–(anti)proton (p–p or p–$\bar{p}$) colliders or electron–positron (e^+–e^-) colliders, and both approaches have significant difficulties in extension to higher energies. Hadrons are composite objects; so only a small fraction of the total energy participates in a collision, and this fraction decreases with increasing energy. Also, production of new particle states is masked by a large background of nonresonant hadronic events; identification of new physics becomes increasingly difficult with increasing energy. Leptonic (e^+–e^-) colliders have had the advantage of providing simple, single-particle interactions with little background. However, extension to higher energies is limited in energy, luminosity, and resolution by radiative effects (synchrotron radiation in circular colliders, beam–beam radiation and pair creation in linear colliders). At very high energies, the collisions are no longer point-like, because of the radiative background.

However, this radiation scales inversely as the fourth power of the lepton mass. Thus, we can extend the high-quality features of e^+–e^- colliders to much higher energies by colliding higher-mass leptons such as muons. The μ^+–μ^- collider concept has been suggested, and is described in some detail in Refs. [1–4], and an example is displayed graphically in Fig. 1. In those initial concepts, a high-intensity multi-GeV hadron accelerator beam produces pions from a hadronic target, and muons are obtained from π-decay. The μ^+'s and μ^-'s are accumulated and cooled by ionization cooling, and then accelerated (in linacs and/or synchrotrons) to high energies for high-energy collisions in a storage ring. The process is repeated at a rate matched to the high-energy muon lifetime to obtain potentially high luminosities.

Since the initiation of the muon collider concept, some subsequent developments have increased interest in the possibility of muon colliders, and recent progress in related fields may increase their potential capabilities.

In high-energy physics (HEP), plans for the current generation of high-energy facilities are now reasonably well established. The next HEP devices are to be high-energy (8–20 TeV) p–p colliders (SSC and/or LHC), to be followed by an e^+–e^- linear collider (up to 0.5 TeV per beam). It is now possible to begin serious consideration of projects to follow these, such as a μ^+–μ^- collider. The current generation of HEP devices includes large components (injectors, tunnels, storage rings, linacs) and technology (high-gradient linacs, low-β^* optics, large-bandwidth stochastic cooling, etc.) which may be incorporated into a muon collider.

Development of high-intensity accelerator concepts (for kaon factories, or accelerator transmutation of waste, or tritium production, or for μ-catalyzed fusion) can also provide methods and facilities for improved μ production, collection and cooling. Some of the possible extensions are suggested below.

Recent studies of e^+–e^- linear colliders show that that approach appears to be limited in energy or luminosity at the TeV scale by beamstrahlung radiation effects. This calls for a "new paradigm" in extension to higher energies [5], possibly including a μ^-–μ^- collider option. Also, the next frontier in HEP appears to be in understanding the "Higgs sector", the generation of masses. The greater direct coupling of muons to the Higgs sector (by $(m_\mu/m_e)^2$) may provide an important incentive for devel-

☆ Work supported by Department of Energy contract #DE-AC05-84ER40150.

SSDI 0168-9002(94)00523-A

oping μ^+–μ^- colliders, possibly in the energy-specific form of a Higgs-resonance "factory".

In this paper we present an overview of possible high-energy μ^+–μ^- colliders for the μ^+–μ^- workshop (Napa, CA, 12/1992), following the concepts previously presented in the Refs. [1–4]. In this overview, we will identify some of the critical problems in implementing the scenario, suggest some possible variations and improvements, and, hopefully, inspire some further directions for innovation and invention from the participants and other readers.

2. μ production

A critical problem in μ^+–μ^- colliders is the production and collection of adequate numbers of muons. The critical difficulties are the large initial phase-space volume of the muons (since they are produced as secondary or tertiary products of high-energy collisions) and the fact that muons decay, with a lifetime of 2.2×10^{-6} γ decay time. (γ is the usual kinetic factor (E_μ/m_μ).) The problem is to obtain, compress, and use sufficient muons before decay.

We will first consider as a baseline reference a high-energy hadronic (HEH) muon source of the type described in Refs. [1,2,4]: a medium-energy high-intensity hadron accelerator beam is transported onto a high-density target to produce GeV-energy π's, and the π's are confined in a transport channel where they decay to produce μ's, which are collected, compressed, and accelerated. In this section, the baseline HEH source is described, and guidelines for optimization and improvement are suggested. Variations and alternative approaches are then mentioned; these include a low-energy (GeV-scale) "π-factory" beam, possibly producing surface muon beams, or an e^- beam producing μ^+–μ^- pairs by photoproduction.

Production of large numbers of muons in the baseline scenario is not difficult. Hadronic interactions of the beam with the target produce large numbers of π's, and almost all of these π's decay by producing a muon plus a neutrino. The difficult problem of optimizing production and collection for maximal μ intensity is not yet solved; however, some guidelines may be obtained from approximate calculations. A first estimate of $\pi^\pm$ production in proton–hadron collisions may be obtained using known and calculated particle spectra, such as the empirical formulae of Wang [6]:

$$\frac{d^2N}{dP\,d\Omega} \cong AP_p X(1-X)e^{-BX^C-DP_t}\frac{\text{pions}}{\text{sr-GeV}/c}\ \text{/interacting proton} \qquad (1)$$

where P_p is the incident proton momentum, $X = P_\pi/P_p$ is the pion/proton momentum ratio, P_t is the pion transverse momentum and $A = 2.385$ (1.572), $B = 3.558$ (5.732), $C = 1.333$ (1.333), and $D = 4.727$ (4.247) for positive (negative) pions. In this formula, pions are produced with a mean transverse momentum of $\sim D^{-1}$ or ~ 0.2 GeV. Also, if $P_\pi/P_p \ll 1$, pion production is nearly independent of proton energy, and pion production within a given momentum bite ($\Delta P_\pi/P_\pi$) is nearly constant. If the Wang model is accurate, an optimal π-source may be obtained from a medium-energy proton beam (20–50 GeV) which collides into a nuclear target, followed by strong-focusing optics which collects secondaries (P_t acceptance ~ 0.3 GeV) and a strong-focusing transport line for $\pi \to \mu\nu$ decay. Small spot sizes on the production target and small beam sizes in the transport line are desired to minimize π and μ emittance. High momentum acceptance ($\delta P/P >$ $\pm 10\%$) is desired for maximal production. Economy would favor lower energies.

Fig. 2 displays a possible configuration for μ^+ or μ^-

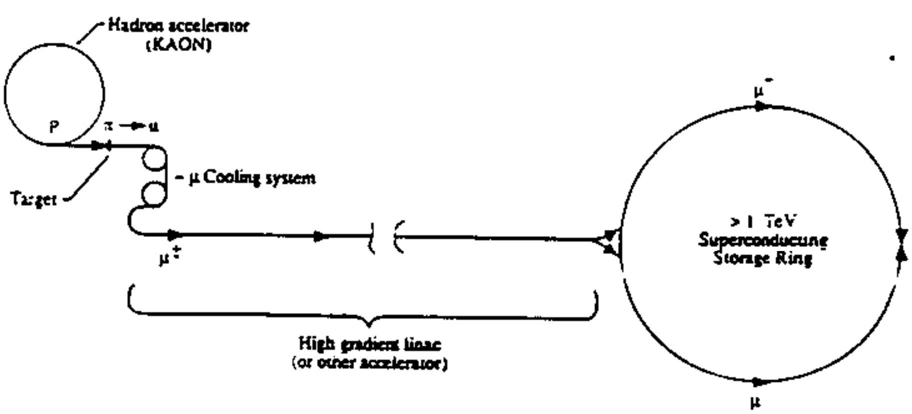

Fig. 1. Overview of linac-based μ^+–μ^- collider, showing a hadronic accelerator, which produces π's on a target, followed by a π-decay channel ($\pi \to \mu\nu$) and μ-cooling system, followed by a μ-accelerating linac, feeding into a high-energy storage ring for μ^+–μ^- collisions. (The linac could be replaced by a rapid-cycling ring; see Fig. 6.)

production. A 40 GeV proton beam is focused (possibly with a Li lens) onto a ~5 cm W target to a mm-scale spot, producing π's, which are captured by the following optics, which is designed to accept 2 GeV π's at angles up to $\Theta_T = 150$ mrad ($P_t < 0.3$ GeV/c). This optics is approximated by a 2.0 cm radius Li lens of 10 cm length centered 15 cm downstream from the target center. The capture optics is to be followed by a strong-focusing transport of ~1 π decay length (~110 m at 2 GeV), which maintains a mean betatron function of 1 m. The resulting muons are inserted into a muon accumulator system for cooling and acceleration.

According to the Wang model, this outline system would produce and accept 0.12 (0.08) π's /(GeV/c)/interacting proton. This must be multiplied by a target efficiency factor $\eta_T \approx 0.4$ (the probability that an incident proton produces an exiting π), and by the momentum acceptance width (0.4 GeV/c for $\pm 10\%$ acceptance), to obtain the number of π's accepted in the decay channel per primary proton (0.019/0.012 $\pi^+\pi^-$).

In the decay channel, π-decay ($\pi \to \mu\nu$) produces μ's within a uniform energy distribution (between 0.57 and $1.0E_\pi$) and with a maximum transverse momentum of 29.8 MeV/c. With 63%(1/e) of the π's decaying and ~40% of the product μ's within the transport acceptance, we find ~0.005 (0.0033) $\mu^+(\mu^-)$ per primary proton are delivered to the μ-collector.

The transverse emittance is determined by the π-production phase space, and the π-decay phase space. The transverse emittance from π-production from a thin target is of order $r_T\Theta_T$ or 3×10^{-4} m rad at 2 GeV, if the primary beam size on the target, r_T, is 0.002 m. The π-decay increases emittance by $\sim \beta\theta_d^2/2$, where β is the mean decay-line betatron function and θ_d is the maximum decay angle. With $\theta_d = P_t/P_\mu = 0.02$ and $\beta = 1$ m in our reference case, an emittance ϵ_t of $\sim 2\times 10^{-4}$ m rad at 1.5 GeV is obtained ($\sim 3\times 10^{-3}$ m rad normalized). (Total emittance from target size, target length, and decay effects should be less than twice this value, and could be reduced somewhat by further optimization.)

This reference case is oversimplified. The problems of separating primary p's from secondary π^+ and π^- beams, and separating the π beams from each other are not addressed; the actual optics will be more interesting. The parameters are not optimized, and the production estimates are not very accurate. The calculations do not include π's (and μ's) from secondary interactions. For $E_\pi \ll E_p$, secondary and cascade production could be large. However, μ-decay from source to collider will reduce the final number (by ~2×). The calculations do show that an HEH source can obtain $\geq 10^{-3}$ μ's per primary proton, and that is adequate for a high-luminosity collider. This results is in agreement with an independent analysis of Noble [7]. It is possible that improvements, acceptance increases, and optimization could increase this to the 10^{-2} level, but probably not much greater.

The HEH scenario has been motivated from a high-energy physics bias, and imitates $\bar{p}$ source methods. π-production also occurs in low-energy hadronic (LEH) "π-factories", from GeV/nucleon protons or deuterons, at a level of ~0.5 π/p. For μ-collider use, a compressor ring would be needed to combine the proton (deuteron) linac beams into sub-μs pulses. The π's are produced at the 100 MeV energy level, where they can be stopped in an absorber to produce 29 MeV/c μ's, which will lose further energy in the absorber. (Stopping times are at ns levels; the μ-lifetime at rest is 2.2 μs.) (Nagamine has proposed extracting such slow muons in a high-intensity surface muon source [8].) An LEH source can produce large numbers of μ's with low energy and momentum spread, which may require little or no further cooling, and the LEH source could be preferable to an HEH source. The difficulty is in extracting sufficient μ's in a small phase-space volume which are suitable for acceleration in a μ^+–μ^- collider. The problem of calculating and optimizing μ production in an LEH configuration is unsolved; it is an important challenge for the reader.

Another possible μ-source can be obtained by colliding multi-GeV e^- beams into a hadronic target; bremsstrahlung would produce μ^+–μ^- pairs by photoproduction, but not as frequently as e^+–e^- pairs. Obtaining the μ's through pair production avoids the phase-space dilution of π-decay; however, μ-production does not seem to be as copious as in a hadronic source. An optimized calculated comparison has not yet been made.

The best possible μ-source is not yet identified or developed. It may follow some of the ideas suggested in this section, with added improvements and innovations, or it may be dramatically different. This is an important challenge for the workshop.

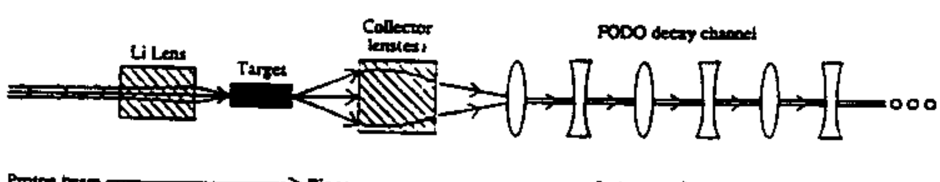

Fig. 2. Schematic view of μ-production from π-decay, with π's produced from hadrons. A high-energy hadronic beam is focused onto a target; a collector lens(es) collects the resulting π's into a strong-focusing "FODO" channel, where π-decay produces μ's for collider use.

3. μ-cooling, and combination

The μ's are produced in a relatively large phase-space volume which must be compressed to obtain high-luminosity collisions. Most of the needed compression is obtained from adiabatic damping; acceleration from GeV-scale μ collection to TeV-scale collisions reduces phase-space by $\sim 10^9$ (10^3 per dimension). Additional phase-space reduction can be obtained by "ionization cooling" of muons ("μ cooling"), which is described in some detail in Refs. [1–3], and is conceptually similar to radiation damping. In this section we first describe transverse μ cooling. Longitudinal cooling and bunch combination, and muon survival and acceleration are then discussed.

The basic mechanism of transverse μ cooling is quite simple, and is shown graphically in Fig. 3. Muons passing through a material medium lose energy (and momentum) through ionization interactions. The losses are parallel to the particle motion, and therefore include transverse and longitudinal momentum losses; the transverse energy losses reduce (normalized) emittance. Reacceleration of the beam (in rf cavities) restores only longitudinal energy. The combined process of ionization energy loss plus rf reacceleration reduces transverse momentum and hence reduces transverse emittance. However, the random process of multiple scattering in the material medium increases the emittance.

The equation for transverse cooling can be written in a differential-equation form as:

$$\frac{\mathrm{d}\epsilon_\perp}{\mathrm{d}z} = -\frac{\frac{\mathrm{d}E_\mu}{\mathrm{d}x}}{E_\mu}\epsilon_\perp + \frac{\beta_0}{2}\frac{\mathrm{d}\langle\theta_{\mathrm{rms}}^2\rangle}{\mathrm{d}z}, \tag{2}$$

where $\epsilon_\perp$ is the (unnormalized) transverse emittance, $\mathrm{d}E_\mu/\mathrm{d}z$ is the absorber energy loss per cooler transport length z, β_0 is the betatron function in the absorber and Θ_{rms} is the mean accumulated multiple scattering angle in the absorber. Note that $\mathrm{d}E_\mu/\mathrm{d}z = f_A \mathrm{d}E_\mu/\mathrm{d}s$, where f_A is the fraction of the transport length occupied by the absorber, which has an energy absorption coefficient of $\mathrm{d}E_\mu/\mathrm{d}s$. Also the multiple scattering can be estimated from:

$$\frac{\mathrm{d}\langle\Theta_{\mathrm{rms}}^2\rangle}{\mathrm{d}z} \cong \frac{f_A}{L_R}\left(\frac{0.014}{E_\mu}\right)^2, \tag{3}$$

where L_R is the material radiation length and E_μ is in GeV. (The differential-equation form assumes the cooling system is formed from small alternating absorber and reaccelerator sections; a similar difference equation would be appropriate if individual sections are long.)

If the parameters are constant, Eqs. (2) and (3) may be combined to find a minimum cooled (unnormalized) emittance of

$$\epsilon_\perp \to \frac{(0.014)^2}{2E_\mu}\frac{\beta_0}{L_R\frac{\mathrm{d}E_\mu}{\mathrm{d}z}} \tag{4}$$

or, when normalized

$$\epsilon_N = \epsilon_\perp\gamma \to \frac{(0.014)^2}{2m_\mu c^2}\frac{\beta_0}{L_R\frac{\mathrm{d}E_\mu}{\mathrm{d}z}} \tag{5}$$

(all energies are in GeV).

Avoiding longitudinal phase-space dilution implies cooling at $E_\mu \geq 0.3$ GeV, and economy implies cooling at relatively low energies (since cooling by e^{-1} requires E_μ energy loss and recovery). $E_\mu \approx 0.5$–1.5 GeV seems reasonable. Cooling can be obtained in either linacs (possibly using recirculation lines) or storage rings. Multiple stages

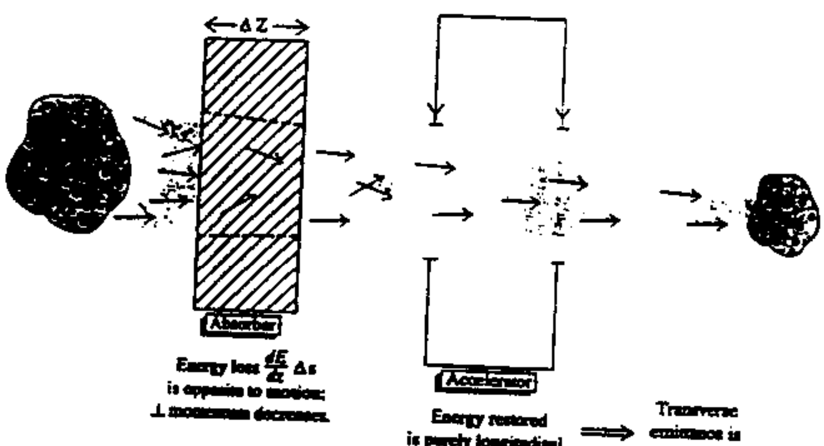

Fig. 3. Schematic view of transverse "ionization cooling". Energy loss in an absorber occurs parallel to the motion; therefore transverse momentum is lost with the longitudinal energy loss. Energy gain is longitudinal only; the net result is a decrease in transverse phase-space area.

can be used to optimize cooling scenarios. The important constraint is that cooling must be completed within a muon lifetime, which can be expressed as $\sim 300\bar{B}$ (T) turns in a storage ring, where $\bar{B}$ is the mean bending field, or as a length $L_\mu = 660\gamma$ m of path length. This constraint may be surmountable.

Some guidelines for optimal cooling may be obtained from Eqs. (2)–(5). They indicate that it is desirable to obtain small β_0 (strong focusing) at the absorber. It is also desirable to have materials with large values of the product $L_R\, \mathrm{d}E/\mathrm{d}s$. $L_R\, \mathrm{d}E/\mathrm{d}s$ is largest for light elements (0.1 GeV for Li, Be but ~ 0.01 GeV for W, Pb), indicating the desirability of light absorbers. However, the need for small β_0 and the depth of focus constraint that a (non-focusing) absorber section must be less than $2\beta_0$ would favor large $\mathrm{d}E/\mathrm{d}s$ (heavy) absorbers. A conducting light-metal absorber (Li, Be, Al) can also be a continuously focusing lens, which could then be arbitrarily long while maintaining small β_0. With current technology, lenses maintaining $\beta_0 < 1$ cm appear possible.

With $\beta_0 = 1$ cm, and $L_R\, \mathrm{d}E/\mathrm{d}s \approx 0.1$ GeV, a normalized emittance of $\epsilon_N \approx 10^{-2}\beta_0 = 10^{-4}$ m rad is obtained as a reasonable goal for transverse cooling. Some improvements may be possible; the reader may develop ideas for optimal implementation.

Longitudinal (energy-spread) cooling is also possible, if the energy loss increases with increasing energy. The energy loss function for muons, $\mathrm{d}E/\mathrm{d}s$, is rapidly decreasing (heating) with energy for $E_\mu < 0.3$ GeV, but is slightly increasing (cooling) for $E_\mu > 0.3$ GeV. This natural dependence can be enhanced by placing a wedge-shaped absorber at a "non-zero dispersion" region where position is energy-dependent (see Fig. 4). Longitudinal cooling is limited by statistical fluctuations in the number and energy of muon–atom interactions. An equation for energy cooling is:

$$\frac{\mathrm{d}\langle(\Delta E)^2\rangle}{\mathrm{d}z} \approx -2\frac{\partial \frac{\mathrm{d}E_\mu}{\mathrm{d}z}}{\partial E_\mu}\langle(\Delta E)^2\rangle + \frac{\mathrm{d}\Delta E_{\mathrm{rms}}^2}{\mathrm{d}s}, \qquad (6)$$

where the derivative with energy combines natural energy dependence with dispersion-enhanced dependence. An expression for this enhanced cooling derivative is:

$$\frac{\partial \frac{\mathrm{d}E_\mu}{\mathrm{d}z}}{\partial E_\mu} = f_\Lambda \left(\frac{\partial \frac{\mathrm{d}E_\mu}{\mathrm{d}s}}{\partial E_\mu}\right)_0 + f_\Lambda \frac{\mathrm{d}E_\mu}{\mathrm{d}s}\frac{\mathrm{d}\delta}{\mathrm{d}x}\frac{\eta}{E_\mu \delta_0}, \qquad (7)$$

where η is the dispersion at the absorber, and δ and $\mathrm{d}\delta/\mathrm{d}x$ are the thickness and tilt of the absorber. Note that using a wedge absorber for energy cooling will reduce transverse cooling; the sum of transverse and longitudinal cooling rates is invariant. In the long-pathlength Gaussian-distribution limit, the heating term or energy straggling term is given by [9]:

$$\frac{\mathrm{d}(\Delta E_{\mathrm{rms}})^2}{\mathrm{d}s} \cong 4\pi(r_e m_e c^2)^2 N_0 \frac{Z}{A}\rho\gamma^2(1-\beta^2/2),$$

where N_0 is Avogadro's number and ρ is the density. Since this increases as γ^2, cooling at low energies is desired. From balancing heating and cooling terms, we find that cooling of $\Delta E_\mu/E_\mu$ to ≤ 0.02 at $E_\mu = 0.25$ GeV is possible, which adiabatically damps to a 1 GeV energy spread of $\Delta E_\mu/E_\mu < 0.005$. An rf buncher plus compressor arc (or synchrotron oscillations) can use this reduced energy spread to obtain reduced bunch length.

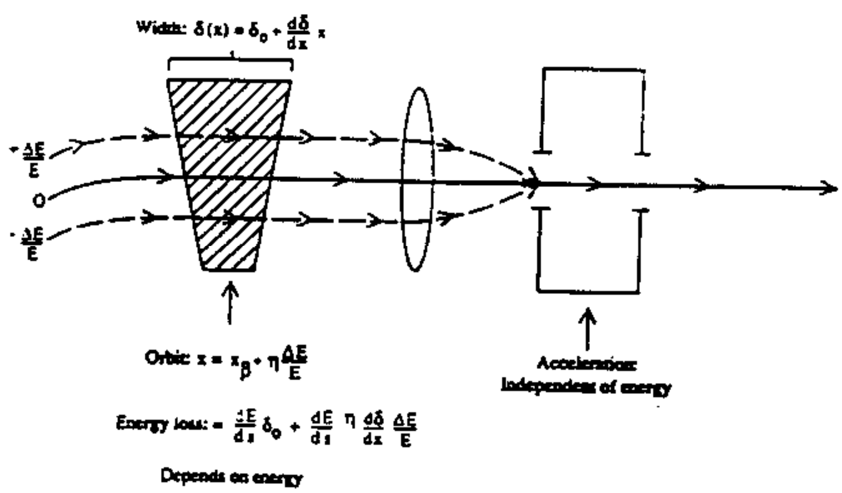

Fig. 4. Enhancement of energy cooling by using a wedge absorber placed in a non-zero dispersion region. The thickness of the absorber depends on transverse position ($\Delta = \delta_0 + (\mathrm{d}\delta/\mathrm{d}x)x$), and the position at the absorber depends on the energy ($x = \eta(\Delta E/E)$), producing enhanced energy dependence of energy loss, decreasing energy spread. Energy recovery in the accelerator is independent of energy. (Transverse cooling decreases with enhanced energy cooling.)

The major problem in longitudinal space derives from the mismatch between the initial bunch structure of the μ-source and the desired μ-collider bunch configuration. The primary proton beam from a rapid-cycling synchrotron (RCS) would consist of ~ 100–200 bunches, while compression to one (or a few) μ-collider bunches is desired. Phase-space manipulations will be needed.

Bunch combination procedures could include:

1) Proton-bunch overlap: Before extraction from the RCS, the proton bunches can be compressed with a lower harmonic (or sideband) rf system to provide spatially overlapping bunches (with different momenta) on the target. Since target spot sizes need not be very small and π-production is not energy-dependent, a broad primary energy spread can be accepted at the target, with no degradation of π-production. Combination by at least a factor of 10 should be obtainable. A separate extraction-energy proton compressor ring for the rf bunch manipulations may be desired, and may be capable of combining the number of bunches to ~ 2–4.

2) Non-Liouvillian "stochastic injection" [10]: Bunch combination without phase-space dilution can occur during π decay, as in "stochastic injection" into a μ-storage ring, as shown in Fig. 5. In this process a train of π bunches from a hadronic target is injected into a decay channel, which is also a zero-dispersion straight-section of a μ-storage ring. The π-bunch spacing is matched to the storage ring period (or a low harmonic). The injected π's are not in the acceptance of the ring; the ring accepts lower-energy particles, in particular, those μ's from π decay in the straight section that are within that acceptance. Successive π-bunch arrivals are timed to overlap an accumulated μ bunch. The π lifetime is ≈ 1% of the μ lifetime, and is naturally matched for decay within the first-turn decay channel, while permitting multiple turn accumulation. At reasonable parameters, μ's from 10–30 π bunches can be accumulated in a single bunch, without large μ-decay losses. (The storage ring can also be used for cooling the accumulated μ's.)

3) Beam cooling with bunch combination: Transverse or energy cooling of μ bunches can compress beams to a degree where bunches can be stacked together, using conventional Liouvillian bunch-combination optics. The stacked bunch can then be further cooled to a phase-space volume ≤ the previous cooled single-bunch size. The process can continue through further stacking and cooling steps. The process adds the complications of multiple beam-combination transport lines and optics; however, combinations of many bunches (10–30) could be obtained.

Some combination of these three (plus other to-be-developed) methods can be used to reduce the number of μ bunches to a few enhanced-intensity bunches. (The first two procedures are complementary: proton bunch stacking naturally combines nearby bunches while stochastic injection more naturally combines widely spaced bunches.) Ionization cooling can then be used for further compression, in bunch length or energy spread, for optimal collider use.

4. Acceleration and collider scenarios

Cooled and compressed muon bunches can be used in high-energy high-luminosity colliders. In this section, we describe some potential scenarios. We use as a reference case a 1 TeV per beam μ^+–μ^- collider (2 TeV in the center-of-mass), as this is the energy scale where the μ^+–μ^- collider may begin to be preferable to e^+–e^- colliders. Table 1 shows a reference case, with relatively conservative choices of parameters.

1) Linac–storage ring. This scenario is displayed graphically in Figure 1, and this is probably the highest-luminosity case. μ^+–μ^- bunches from the collector/cooler are both accelerated to full energy in a high-gradient linac to 1 TeV, where both bunches are injected into a superconducting storage ring for high-energy collisions at low-β_0 interaction points. The μ beam lifetime is ~ $300B$ turns, where B is the mean bending field in T. $B \approx 8$ T, imply-

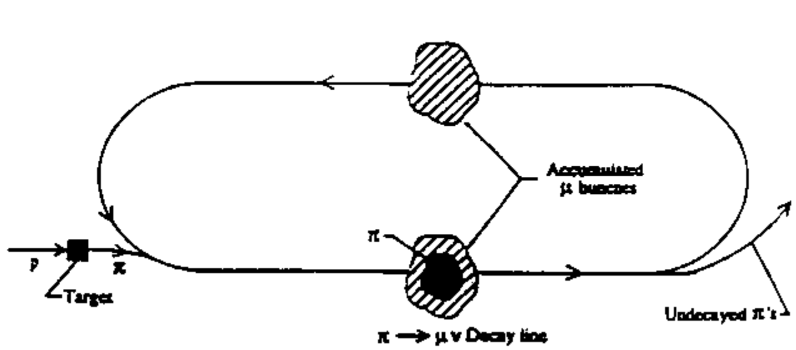

Fig. 5. Schematic view of "stochastic injection" into a storage ring. A train of p bunches produces π bunches, which are injected into a storage ring for multi-turn stacking. The initial spacing is matched to a ring harmonic ($h = 2$ in the figure), so that following π bunches overlap accumulated μ's from previous bunches. π decays in the straight section which produce μ's within the ring acceptance add to the accumulation. (Note that the short but finite π lifetime is nearly optimally matched to make this scheme practical.)

Table 1
Parameter list for TeV μ^-–μ^- collider (2 TeV collisions)

Parameter	Symbol	Value
Energy	$E_{\mu\pm}$	1 TeV
Luminosity	L	3×10^{32} cm^{-2} s^{-1}
HEH-source parameters		
Proton energy	E_p	40 GeV
P/pulse	N_p	10^{14}
Pulse rate	f_0	30 Hz
μ production efficiency	μ/p	10^{-3}
Collider parameters		
#μ^+/μ^- per bunch	$N^{\pm}$	10^{11}
#bunches	n_B	1
Storage turns	n_s	1200
μ emittance	$\epsilon_\perp = \frac{\epsilon_N}{\gamma}$	10^{-8} m rad
Interaction focus	β^*	1 mm
Beam size	σ	3 μm

ing 2400 turns, is currently achievable. The main difficulty is the relatively large cost of the full-energy TeV linac.

2) Linac–linac collider. As in e^+–e^- linear colliders, μ bunches from opposing linacs can collide. This scenario loses the luminosity magnification obtained from multiple collisions in a storage ring, which is permitted by the long μ lifetime. It also requires two full-energy linacs, and a TeV storage ring is cheaper than a TeV linac. However, an existing e^-–e^- linear collider could be modified to obtain μ^+–μ^- collisions, with the addition of a μ source.

3) Rapid-cycling synchrotron collider. At 1 TeV, the μ lifetime has increased to 0.021 s, and the lifetime increase with energy is sufficient to permit acceleration in a rapid-cycling synchrotron with acceptable losses. Fig. 6 shows the basic components: a μ source injected into a $\sim$ 20–50 GeV linac followed by a rapid cycling synchrotron with 20 km circumference ($B \approx 1$ T). Acceleration from injection to full energy in $\sim$50–100 turns could follow a 60–120 Hz waveform followed by $\sim$0.02 s at fixed field for collisions or, for higher luminosity, transfer to a fixed-field (8 T) collider ring. Luminosity would be naively expected to be about an order of magnitude smaller than in the linac–storage ring scenario.

4) The μ–p collider. A μ^+–μ^- collider can also be operated as a μ^-–p collider with both μ and p beams at full energy. High luminosity would be relatively easily obtained because only one beam (μ) is unstable and diffuse. This is a probable initial and debugging operating mode for a storage ring μ^+–μ^- collider. Revolution frequencies of equal energy μ and p beams would be naturally mismatched because of unequal speeds. They can be rematched by displacing the beams in energy and using the ring nonisochronicity [2]. The required energy displacement is

$$\frac{\delta E}{E} = \gamma_T^2\left(\frac{1}{2\gamma_p^2} - \frac{1}{2\gamma_\mu^2}\right) \tag{8}$$

where γ_T^2 is the ring transition-gamma. At reference parameters ($E_\mu = 1$ TeV, $\gamma_T = 30$) $\delta E/E = 0.0005$ is required.

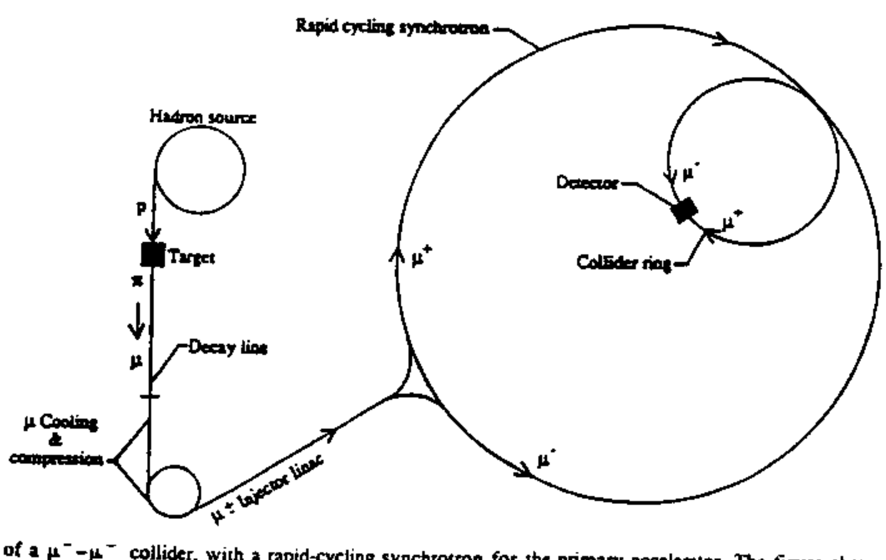

Fig. 6. Overview of a μ^-–μ^- collider, with a rapid-cycling synchrotron for the primary accelerator. The figure shows a primary proton source, producing π's on a target, which decay to μ's in a decay channel. After cooling and compression, μ^+–μ^- bunches are accelerated in a linac and in a rapid-cycling synchrotron to full energy, where they are injected into a high-field storage ring for multi-turn collisions. For example, a 20 GeV linac feeding into a 20 GeV/turn (up to 1.2 T) rapid-cycling synchrotron would produce 1 TeV μ^-–μ^- beams with acceptable decay losses.

5) Physics–opportunity colliders. A resonance, such as a new Z particle or a Higgs particle, may exist or be predicted in μ^+–μ^- collisions. At such a resonance, a lower-luminosity and/or lower-energy collider could still provide extremely important physics. Such a collider would be a simplified form of the baseline high-luminosity models, possibly with μ sources using existing accelerators, omitting μ cooling, and/or using existing storage rings for collisions. Such a facility would, of course, provide excellent training for a high-energy high-luminosity collider.

5. Luminosity possibilities and constraints

Using the previously discussed techniques, high luminosity can be obtained in a μ^+–μ^- collider. The luminosity is given by the equation:

$$L=\frac{f_C N^+N^-}{4\pi\sigma^2}=\frac{f_C N^+N^-}{4\pi\beta^*\epsilon_\perp}. \tag{9}$$

where N^+,N are the number of μ^+, μ^- per colliding bunch, f_C is the frequency of bunch collisions, σ^2 is the colliding beam size, β^* is the betatron function at the collision point and $\epsilon_\perp=\epsilon_N/\gamma$ is the transverse emittance. In a storage-ring collider, $f_C=f_0 n_B n_S$, where f_0 is the system cycling rate, n_B is the number of bunches, and n_S is the number of turns of storage per cycle.

This formula is applied to a reference TeV μ^+–μ^- collider (Table 1). The parameters we use include $N^+=N^-=10^{11}$ (which can be obtained assuming a modest production rate of 10^{-3} μ/p from 10^{14} proton high-energy pulses), $n_B=1$, $f_0=30$ Hz, and $n_S=1200$ turns storage. With $\epsilon_N=10^{-4}$ m rad ($\epsilon_\perp=10^{-8}$) from $\beta_0=1$ cm and $\beta^*=1$ mm ($\sigma\approx 3$ μm), we obtain a respectable baseline luminosity of $L\approx 3\times10^{32}$.

Note that the above parameter set is relatively modest, and improvements in some of the parameters (i.e., N^-, N^-, σ) by up to an order of magnitude are conceivable. However, reliable accomplishment of high luminosity in a novel and complicated facility which uses unstable particles and has several difficult design components will still not be easy.

Luminosity in a μ^+–μ^- collider ring may be expected to be limited by the beam–beam interaction. Since long-term stability is not needed, the allowable beam–beam tune shift should be somewhat greater than the e^+–e^- storage limit of $\Delta\nu\lesssim 0.05$. The tune shift is given by

$$\Delta\nu=\frac{N r_\mu \beta^*}{4\pi\gamma\sigma^2}. \tag{10}$$

where $r_\mu=1.363\times10^{-17}$ m. At the baseline parameters $\Delta\nu\approx 0.001$, luminosity would have to be increased dramatically for the beam–beam limit to be significant.

A μ^+–μ^- collider has substantial beam power requirements, particularly in the primary proton beam. In the reference case, the requirements are 10^{14} 40 GeV protons at 30 Hz, which implies 20 MW beam power. This is an order of magnitude above present facilities, but is the same magnitude as proposed K-factories. (The high-energy μ beams themselves require only 0.32 MW.) More efficient μ-production may be desirable (obtaining more μ/p, using lower-energy p's or a LEH source). However, a high-luminosity high-energy μ^+–μ^- collider is a successor to SSC- or TLC-size facilities, and on that scale, a K-factory type source is small. An even higher intensity source could be affordable.

A significant difficulty in a storage ring is that μ's decay ($\mu\to e\nu\nu$), and decay electrons at ~ 0.3 TeV will hit the walls of the storage ring. At the reference parameters, with half the μ's decaying during storage, we find 50 kW of ~ 0.3 TeV electrons will be deposited evenly within a narrow strip on the inner wall of the ring. The ring must be designed to accept this.

The obtainable luminosity L is expected to increase with increasing end point energy E_μ, as the μ lifetime increases and the emittance and energy spreads are adiabatically damped. As discussed in Ref. (4), the beam size (σ^2) at collision should decrease as E_μ^{-2}, as both β^* and $\epsilon_\perp$ can decrease. Cycle time would increase; however, the longer cycle time could permit accumulation of successive rapid-cycling proton pulses to obtain magnified-intensity μ-bunches. The net effect is that N^+ and N^- increase as E_μ and f_C is reduced by E_μ^{-1}. In all, L naturally increases as E_μ^3. (Costs increase linearly with E_μ.) This scaling would be expected to dominate until the $E_\mu\approx 100$–1000 TeV region, where μ synchrotron radiation excludes μ storage-ring colliders.

6. Summary

In this paper, we have introduced concepts which show the promise of the development of high-luminosity TeV-scale μ^-–μ^- colliders. These initial concepts need considerable practical development. While key ingredients of a future facility have been introduced, further innovations and improvements are greatly desired. These concepts plus further developments must be integrated into a fully self-consistent design for a μ^+–μ^- facility. The discussions at the first μ^+–μ^- collider workshop at Napa, California, should elucidate the possibilities and set a basis for further development. Contributions and improvements from the workshop participants and other readers are encouraged. A functional μ^+–μ^- collider will be obtained only through further innovation and development.

Acknowledgements

I thank D. Cline, A. Ruggiero, J.D. Bjorken, S. Glashow, J. MacLachlan, H.A. Thiessen, R. Palmer and R.

Noble (and any others I forgot) for important conversations contributing to the ideas expressed in this paper.

Appendix

Since this report is appearing subsequent to the 1992 Napa $\mu^+-\mu^-$ Collider workshop, in this section I am adding some initial impressions of the proceedings of the workshop.

D. Cline presented an interesting immediate application for a $\mu^+-\mu^-$ collider [11]. There are theoretical reasons to believe that the Higgs boson might exist in the 90–180 GeV region, and that is an energy region at which only a $\mu^+-\mu^-$ collider could obtain a clean observation of a Higgs boson. ($\mu^+-\mu^- \to H$ is favored because of the relatively large muon mass.) The observation would require luminosity greater than $\sim 10^{29}$ cm^{-2} s^{-1}. Sample parameters with this luminosity goal are shown in Table 2. This relatively low-energy collider could be developed relatively inexpensively, possibly using existing facilities for major components, although it is unclear if any existing accelerator could deliver sufficient muon intensity. An important future goal is developing an optimal short-term path to this extremely important physics goal.

At the workshop, it was speculated that maximal muon production could be obtained from a hadronic (or leptonic) cascade source, possibly at a beam dump, rather than the single-interaction source outlined above. H. Thiessen suggested that the most efficient source could be a ~ 5 GeV hadronic source. Serious target problems were noted for any high-intensity source. Further study/optimization is needed.

In beam cooling, the limitations on ionization cooling due to multiple scattering (described above) were discussed. It may be possible to have a muon source with initial emittance smaller than the multiple scattering limit, and therefore to avoid cooling.

The bunch combination/compression (see above) was identified as a key problem, particularly bunch length reduction to match small β^* optics. At multi-GeV energies, the muons will be relativistic and have no longitudinal motion within a linac. However, relativistic muon bunches can be compressed with an rf-induced energy tilt and transported through a bending arc (or ring). Scenario design/optimization is needed.

Although the workshop did not identify a clear path to a sufficient luminosity design, it did show important (and perhaps obtainable) physics goals, particularly the Higgs discovery opportunity. Future study goals, particularly in source design and scenario development, were identified, and the need for future workshop(s) after further development was suggested.

Table 2
Parameter list for 100 GeV $\mu^+-\mu^-$ collider

Parameter	Symbol	Value
Energy	$E_{\mu\pm}$	100 GeV
Luminosity	L	10^{29} cm^{-2} s^{-1}
Pulse rate	f_0	10 Hz
Storage turns	n_s	1000
#bunches	n_B	10
#μ^+/μ^- per bunch	$N^\pm$	10^{10}
μ-emittance	$\epsilon_\perp = \frac{\epsilon_N}{\gamma}$	10^{-7} m rad
Interaction β	β^*	1 cm
Beam size (at IR)	σ	30 μm

References

[1] D. Neuffer, Particle Accelerators 14 (1983) 75.
[2] D. Neuffer, Proc. 12th Int. Conf. on High Energy Accelerators, eds. F.T. Cole and R. Donaldson ibid., 1983, p. 481.
[3] E.A. Perevedentsev and A.N. Skrinsky, ibid., 1983, p. 485.
[4] D. Neuffer, in: Advanced Accelerator Concepts, AIP Conf. Proc. 156 (1987) 201.
[5] M. Tigner, Proc. 3rd Int. Workshop on Advanced Accelerator Concepts, Port Jefferson, NY, AIP Conf. Proc. 279 (1993) p. 1.
[6] C.L. Wang, Phys. Rev. D 9 (1973) 2609 and Phys. Rev. D 10 (1974) 3876.
[7] R. Noble, Proc. 3rd Int. Workshop on Advanced Accelerator Concepts, Port Jefferson, NY, AIP Conf. Proc. 279 (1993) p. 949.
[8] K. Nagamine, unpublished communication (1986).
[9] U. Fano, Ann. Rev. Nucl. Sci. 13 (1963) 1.
[10] D. Neuffer, IEEE Trans. Nucl. Sci. NS-28, (1981) 2034.
[11] D. Cline, unpublished communication (1992).

Nuclear Instruments and Methods in Physics Research A 350 (1994) 36–44
North-Holland

NUCLEAR INSTRUMENTS & METHODS IN PHYSICS RESEARCH
Section A

Characteristics of a high energy $\mu^+\mu^-$ collider based on electro-production of muons

William A. Barletta *,+, Andrew M. Sessler ++
Lawrence Berkeley Laboratory, Berkeley, CA, USA

Received 8 April 1994

We analyze the design of a high energy $\mu^+\mu^-$ collider based on electro-production of muons. We derive an expression for the luminosity in terms of analytic formulae for the electron-to-muon conversion efficiency and the electron beam power on the production target. On the basis of studies of self-consistent sets of collider parameters under "realistic" ("optimistic") assumptions about available technology with beam cooling, we find the luminosity limited to 10^{27} cm^{-2} s^{-1} (10^{28} cm^{-2} s^{-1}). We also identify major technological innovations that will be required before $\mu^+\mu^-$ colliders can offer sufficient luminosity (10^{30} cm^{-2} s^{-1}) for high energy physics research.

1. Introduction

Many physicists consider that the recent determinations of lower bounds for the mass of the top meson reinforce arguments that a Standard Model Higgs should have a mass less than twice the mass of the Z. This consideration has led to renewed interest in muon colliders as an ideal means of probing the mass range from m_Z to $2m_Z$. More generally, a muon collider with center-of-mass energy in the range of 200 to 400 GeV has the potential to produce very large numbers of Higgs particles because of the enhanced (vis a vis électrons) muon coupling to the Higgs. For such a collider to have maximum discovery potential the luminosity should be $\geq 10^{30}$ cm^{-2} s^{-1} [1]. As the muon is an unstable particle, the muons must be generated as secondary beams from either a proton beam or an electron beam striking a production target. The muons that emerge from the target must then be gathered and accelerated rapidly to high energy, at which point they can be injected into a storage ring collider with superconducting magnets.

This paper analyses the possibility of using electro-production to generate the muon beams. The chief advantage of producing the muons with an electron beam from a high energy, linear accelerator is that the bunches of muons are naturally formed with a short bunch length (< 1 cm) for acceleration to the desired high energy in a subsequent linear accelerator. As the muons will retain their short bunch length in the collider, a low β interaction region can be employed. This scheme is illustrated in Fig. 1.

* Corresponding author.
+ Work performed under the auspices of the Lawrence Livermore National Laboratory under contract W-7405-eng-48.
++ Work performed under the auspices of the Lawrence Berkeley Laboratory under contract DE-AC03-76SF00098.

2. Electro-production

Muons can be produced by an electron beam via two classes of processes, 1) $\mu^+\mu^-$ pair production and 2) photo-production of π's and K's, which subsequently decay into muons. It is known experimentally [2] that the cross-section for pair production is much more than an order of magnitude greater than that for process (2). Consequently, in the discussion that follows we will consider only pair production.

To estimate the muon pair production from an electron beam of energy, E_e, incident upon a thick target of atomic number Z, one can use the expression from Nelson [3] based on approximation A of shower theory. F is number of muons per electron produced at an angle $\geq \phi$ with respect to the incident electron beam:

$$\frac{\mathrm{d}F}{\mathrm{d}E}(E_e, E, \phi) = \frac{1}{(2\pi)^2}\frac{m^2}{\mu^2}\frac{0.572 E_e \eta}{\mu^2 \ln(183 Z^{-1/3})} 2\ln(\gamma_\mu) \times\left\{(1-\nu^2) - 0.33\left[1 - 4\nu^2(1-0.75\nu)\right] \times\eta\left[1+\gamma^2\right]\right\}, \quad (1)$$

where m = electron mass, μ = muon mass, E = energy of

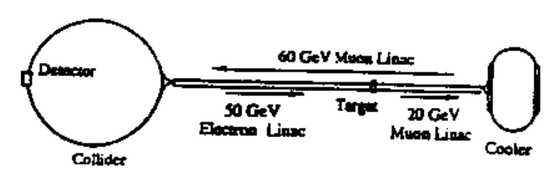

Fig. 1. The scheme for a μ^+/μ^- collider using electro-production.

SSDI 0168-9002(94)00571-8

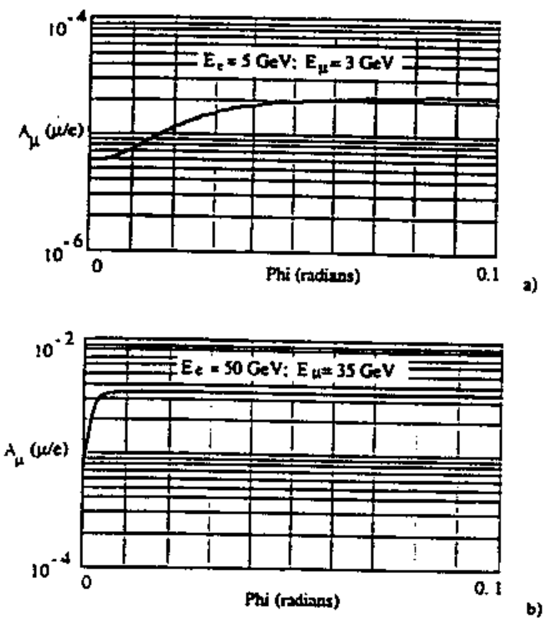

Fig. 2. Number of μ pairs per e^- accepted at an angle $\leq \phi$ for (a) a 5 GeV electron beam with $E_\mu = 3$ GeV and (b) a 50 GeV electron beam with $E_\mu = 35$ GeV.

the muon at the production target, $\gamma_\mu = E/\mu$ at the production target, $\nu = E/E_e$, $\lambda = \gamma^2\phi^2$, and $\eta = (1+\lambda)^{-2}$. Eq. (1) is known to overestimate the muon pair production by a factor of two.

The number of muons per electron accepted in an angle $\leq \phi$, in a momentum bite of $\pm\Delta p/p$ at a muon energy E is

$$A_\mu = \left[\frac{\mathrm{d}F}{\mathrm{d}E}(E_e, E, 0) - \frac{\mathrm{d}F}{\mathrm{d}E}(E_e, E, \phi)\right] E\frac{2\Delta p}{p}. \quad (2)$$

From Eq. (2) it is immediately obvious that one will prefer to accept muons with a large value of E/E_e rather than with a small E/E_e as long as the function $\mathrm{d}F/\mathrm{d}E$ is relatively flat in energy. For small muon production angles this condition obtains for the energy range, $0.2 < E/E_e <$ 0.8. The same consideration also argues that one should choose a large initial electron beam energy. Figs. 2a and 2b display plots of Eq. (2) for a low energy and a high energy production option respectively.

At the front surface of the production target the electron beam can be focused to a spot of radius, $r_b \approx 1$ mm. The muons will, however, appear to originate from a somewhat larger spot with a size given by the radial extent of the electromagnetic shower at a depth corresponding to the shower maximum, which occurs approximately six radiation lengths ($6X_0$) inside the target. The radiation length, X_0, for tungsten is 3 mm; hence the shower maximum will occur at ≈ -20 mm, and a tungsten target 30 mm long will yield almost the entire thick target conversion to muons.

At the depth corresponding to the shower maximum the primary electron beam will have suffered a mean scattering angle of

$$\Theta^2 = \left(\frac{0.028\text{ GeV}}{E_e}\right)^2\left(\frac{6X_0}{X_0}\right), \quad (3)$$

which will induce a radial spread of $6X_0\Theta$ in the primary beam. Actually in a high Z target, the shower will spread by an amount roughly double this value. Hence, the shower radius can be approximated by

$$r_{sh} = \left(r_b^2 + (12\Theta X_0)^2\right)^{1/2}. \quad (4)$$

At production the geometrical emittance, $\epsilon(E)$, of the muon beam of energy, E, accepted into an angle ϕ_{accept} will be

$$\epsilon(E) = \frac{\epsilon_{n,prod}}{\gamma_{prod}} = r_{sh}\phi_{accept}, \quad (5)$$

where $\epsilon_{n,prod}$ is the normalized emittance at production.

To increase the muon production efficiency one might consider alternate techniques of photo-production. The production process consists of two steps: 1) conversion of the electron energy into photons and 2) muon pair production from the photons. Rather than using bremsstrahlung, one might employ synchrotron radiation as the conversion process. Synchrotron radiation conversion could either take place in a crystal or in a plasma [4], which has an obvious advantage of being more amenable to high average power operation.

The choice of synchrotron radiation conversion is unlikely to increase the rate of muon production as the mean photon energy is lower for the synchrotron radiation photons than for the bremsstrahlung photons. The synchrotron radiation photons are more numerous, but only at low energies for which muon pair production is not energetically allowed. The angular distribution of the muons produced will be dominated by the spread of angles of the electrons in the primary beam as the average electron angle will be significantly larger than γ^{-1}.

The pair production rate in crystals is known experimentally [5] to be larger than in amorphous materials due to the coherent field effects. For photons of 100 GeV, the coherent production is a few times the Bethe–Heitler rate; however, for 20 GeV photons this effect increases pair production by only 10%. As the mean energy of bremsstrahlung photons is $\approx 20\%$ of the incident beam energy, pair production in a crystal will not significantly enhance the muon yield for a 100 GeV per beam collider. Hence, in the analysis that follows we restrict our attention to the use of a conversion bremsstrahlung production target.

3. Ionization cooling

In designing a collider one will inevitably seek a means of having as low an emittance as possible for the beams. One suggested means of cooling the muons (Fig. 3) is to pass the beam through a succession of alternating slabs of material (ionization cells) and rf-accelerating sections. In the ionization cells each of the muons gives up momentum along its particular trajectory, thereby losing transverse and longitudinal momentum. In the accelerating sections the longitudinal momentum is restored to the beam. Thus the transverse emittance of the beam is reduced in a manner analogous to radiation damping.

Neuffer [6] has shown that the ionization cooling of the transverse emittance is limited by beam heating due to multiple Coulomb scattering. If the transverse cooling is performed at an energy, E_c, using a medium for which the radiation length is X_R and the ionization loss rate is dE/dx, then the equilibrium, normalized emittance will be

$$\epsilon_{eq,n} = \frac{\beta_{cool}}{2} \frac{(14\ \mathrm{MeV})^2}{m_\mu c^2 \left(X_R \dfrac{dE}{dx} \right)}, \tag{6}$$

where β_{cool} is the value of the beta function in the scattering medium. From Eq. (6) it follows that efficient cooling requires that one employ a very strong focusing system that brings the beam to a symmetric waste of small radius in the ionization medium. For high energy muons traversing a medium of density ρ (g/cm^3), of atomic number Z, and of atomic weight, A, the ionization loss rate can be approximated [7] by

$$\frac{dE}{dx} = \frac{DZ\rho}{A\beta^2} \left\{ \ln\left(\frac{2m_e\gamma^2\beta^2c^2}{I} - \beta^2 \right) \right\}, \tag{7}$$

where $\beta = v/c$, $D = 0.307$ and $I = 16Z^{0.9}$ eV. For materials with $Z \geq 6$, the radiation length may be approximated by

$$X_R(\mathrm{cm}) = \frac{716.4A}{\rho Z(Z+1)\ln(287Z^{-1/2})}. \tag{8}$$

Multiplying Eqs. (7) and (8), one observes that the product $(X_R\ dE/dx)$ is independent of the density of the ionization medium and is greatest for small values of Z. Hence, low Z media will be preferred over high Z media for ionization cooling. The length of the scattering medium in any individual ionization cell will have to be limited to $\beta_{cool}/2$.

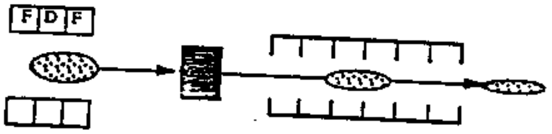

Fig. 3. Schematic of the basic components of an ionization cooling array: a strong lens to focus the beam, the ionization medium in which the particles lose both transverse and longitudinal momentum, and an accelerating structure to restore the longitudinal momentum of the beam.

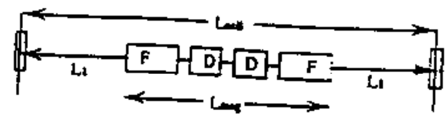

Fig. 4. Schematic of the triplet optics of an ionization cooling cell; the disks of the ionizing medium are shaded and have a half width of $\beta_{cool}/4$.

As the momentum bite of the selected muons will be relatively large, one should consider using optics with second order chromatic corrections to focus the beam onto the ionization targets as otherwise the spot size will be unacceptably large. The focusing system may be a strong quadrupole triplet. Brown [8] has suggested a focal system that is suitable for scaling calculations. In this triplet transverse dimensions are scaled by a factor a_q, which is the aperture (radius) of the first quadrupole of the triplet; longitudinal dimensions are scaled by the "ideal" focal length, f,

$$f = \left(\frac{a_q}{B_q}(B\rho) \right)^{1/2}, \tag{9}$$

where B_q is the pole tip field in the first quadrupole, and $(B\rho)$ is the magnetic rigidity of the beam. For a beam of momentum p,

$$B(\mathrm{T})\rho(\mathrm{m}) = 3.3p(\mathrm{GeV}/c)c. \tag{10}$$

The optical invariants of this particular triplet design are incorporated into the scaling equations that follow; the geometry of the design is illustrated in Fig. 4.

The free space from the focus to the first quadrupole, L_t, is $1.36f$; the length of the triplet, L_{trig}, is $3.13f$ and the length of the cooling cell, l_{cell} is $5.85f$. Without chromatic correction the value of β_{cool} for a beam with fractional momentum spread $\sigma_p (= \Delta p/p)$ is given by

$$\beta_{cool} = 5.92\sigma_p f. \tag{11}$$

With second order chromatic correction of the focusing optics the beta function can be reduced to

$$\beta_{cool} = 74.0(B\rho)\left(\frac{a_q}{fB_q} \right)\sigma_p^2. \tag{12}$$

In the analysis that follows we chose the corrected optics described by Eq. (12). As the cooling disks have a length, $\beta_{cool}/2$, each cooling cell produces an energy loss of e_{cell}, limited to

$$e_{cell} = \frac{\beta_{cool}}{2}\frac{dE}{dx}. \tag{13}$$

In optimizing the production/cooling scenario for the muon collider, one can now choose both the energy of muon production, E_μ, and the energy at which the cooling is performed, E_c. Note that although the equilibrium emittance of Eq. (6) does not depend explicitly on the cooling energy, the choice of E_μ, E_c, and the momentum acceptance will determine σ_p and thereby β_{cool} in the cooling lattice. Thus the choice of E_c will determine the transverse cooling coefficient, C_μ, via

$$C_\mu = \frac{\epsilon_{n,prod}}{\epsilon_{eq,n}} = \frac{r_{sh}\phi_{accept}\gamma_{prod}}{\epsilon_{eq,n}}. \qquad (14)$$

The choice of E_μ will also influence the number of muons per bunch that are available to be injected into the collider as some of the muons will decay as they traverse the cooling lattice. If energy is replaced during the cooling process by accelerator cells with an average accelerating gradient G, the total path length in the cooling lattice, L_{cool}, will be

$$L_{cool} = \frac{E_c}{F_c}\left(\frac{L_{cell}}{e_{cool}} + \frac{1}{G}\right)\ln(C_\mu). \qquad (15)$$

In Eq. (15) F_c is the overall packing fraction of the ionization and acceleration cells in the cooler lattice. F_c accounts for pumping ports, flanges, diagnostics, bending magnets, and sextupoles in the cooling lattice. If the number of muons per bunch that is injected into the cooling lattice is N_μ, then the number of muons per bunch available for injection into the collider will be

$$N_\mu^* = N_\mu \exp\left(-\frac{L_{cool}}{c\tau_\mu\gamma_c}\right), \qquad (16)$$

where τ_μ is the muon lifetime at rest, and γ_c is $E_c/m_\mu c^2$.

Longitudinal cooling of the beam would allow smaller values of β_{cool} and consequently lower equilibrium emittances. Such reduction of the momentum spread can be accomplished by two means: 1) adiabatic damping by accelerating the muons prior to transverse cooling and 2) ionization cooling either in the transverse damper or in a separate damping structure. If the longitudinal cooling were limited to the ionization damping in the zero-dispersion cells of the transverse damper described above, the amount of acceleration, A_μ needed [4] to reduce the momentum spread by a factor $1/e$ would be

$$A_\mu = E_c\left(\frac{\mathrm{d}E/\mathrm{d}x}{E_c\dfrac{\partial^2 E}{\partial E_\mu\,\partial x}}\right) \approx 5E_c. \qquad (17)$$

If the longitudinal cooling is done in a dispersive section, the energy spread might be reduced by $1/e$ with as little as $2E_c$ of total energy exchange.

As computed from Eq. (16), the path length of the muons in the cooler will typically be tens of kilometers, even if the packing fraction of the cooling lattice is large. A large packing fraction in conjunction with a high muon

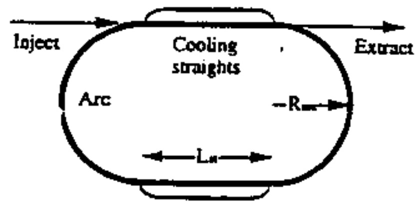

Fig. 5. A 2-in-1 muon cooling ring. Transverse coolers are in each of the straight sections. The gray sections have large aperture quadrupoles for the first stage of cooling; the black straights have stronger, small aperture quadrupoles.

energy implies that the transverse emittance cooler should be constructed in the form of a recirculating linac such as CEBAF with high field bending magnets in the arcs and with as much as a few GeV per turn of acceleration in the straight, cooling sections.

At injection into the cooler the transverse emittance and the momentum spread of the muon beam will be large. Consequently the apertures of the quadrupoles in the cooling straights must be relatively large. One may envision a more effective form of cooler in which the emittance is reduced by an order of magnitude before injection into a final cooler which can have stronger, smaller aperture quadrupoles. As the value of β_{cool} can be much smaller in the second cooler, the equilibrium emittance could be much smaller than achievable in a single cooling ring. In such a scheme there is no need to duplicate the cost of the high field, dipole arcs. Rather the two coolers can share common arcs in a 2-in-1 arrangement illustrated in Fig. 5. The choice of straight-through or by-pass paths for the cooling stages can be selected to minimize the total path length of the muons in the coolers.

In the cooling ring the total length of the cooling cells plus re-acceleration cavities is $2L_{st}P_c$ where P_c is the packing fraction of ionization cells plus accelerator cells in the straight sections. As the overall packing fraction, F_c, is just $[2L_{st}P_c(2\pi\ R_{arc} + 2L_{st})^{-1}]$, the number of cooling cells, N_c is related to the average dipole field in the bends, $\langle B_d\rangle$ and the accelerating field, G, by

$$N_c = F_c\frac{2\pi(B\rho)}{\langle B_d\rangle}\left[1 - \frac{F_c}{P_c}\right]^{-1}\left[l_{cell} + \frac{e_{cell}}{G}\right]^{-1}. \qquad (18)$$

Hence, the rf-system of the cooling ring must supply $N_c e_{cell}$ volts per turn. In damping the emittance of the muons by a factor C_μ, the muons must execute $[E_c N_c e_{cell}^{-1} \ln C_\mu]$ turns.

4. Collider considerations

The number of muons per bunch, N_μ^*, that circulate in the collider will be determined by the production efficiency, A_μ, by the charge, N_e in the electron bunch that strikes the production target, and by the path length through the cooling lattice. The number of electrons per bunch will be limited by the beam loading in the linac and by the design of the electron gun. The present SLAC gun (thermionic) produces bunches of 10 nC. If the electron beam emittance is not critical, the charge in the electron bunch can be raised to 20–30 nC. Bunches with as much as 50 nC may be produced with a photocathode gun, but such a large charge would lead to large beam loading and complications form the beam-breakup instability in an S-band linear accelerator.

The electron bunches will be produced in a macropulse of duration, τ_e, which is chosen to match the circulation period of the muons in the storage ring collider (Fig. 1). If the average dipole field in the ring is 3 T and if the muon energy is 100 GeV, then the circulation period will be 2 μs:

$$\tau_e = 2\mu\text{s} \times \left(\frac{3\ \text{T}}{B_{\text{ave}}}\right)\left(\frac{E_\mu}{100\ \text{GeV}}\right). \tag{19}$$

If the number of bunches per macropulse is N_b, then the frequency of collisions in the collider will be

$$f_{\text{coll}} = N_b/\tau_e. \tag{20}$$

To maintain the muon population in the collider the linac must be pulsed at a frequency of τ_μ^{-1}, where τ_μ is the muon lifetime as seen in the laboratory; at 100 GeV, $\tau_\mu = 2$ ms. Hence, the duty factor of the linac will be τ_e/τ_μ. The average power of the electron beam on the muon production target is, therefore,

$$P_{\text{beam}} = q\frac{N_b N_e}{\tau_e}\frac{\tau_e}{\tau_\mu}E_e = q\frac{N_b N_e}{\tau_\mu}E_e, \tag{21}$$

where q is the electron charge.

The peak luminosity of the collider with muons with a geometrical emittance, ϵ, can be written as

$$L = \frac{N_\mu^{*} f_{\text{coll}}}{4\pi\epsilon\beta^*}, \tag{22}$$

where $\gamma = E_\mu/\mu$ and β^* is the value of the beta function at the collision point. Combining Eqs. (5), (16), (20) and (22), and evaluating the average luminosity of a collider of repetition rate, R, we obtain the following expression for the average luminosity of the collider,

$$\langle L\rangle = \frac{A_\mu^2 N_e^2 N_b \gamma C_\mu}{4\pi r_{\text{sh}}\phi_{\text{accept}}\beta^*\tau_e}\left(\frac{\gamma}{\gamma_{\text{prod}}}\right)\exp\left(-\frac{2L_{\text{cool}}}{c\tau_\mu\gamma_c}\right) \times\left[1-\exp\left(\frac{2}{\tau_\mu\gamma R}\right)\right]\left(\frac{\tau_\mu\gamma R}{2}\right), \tag{21}$$

The factor, $\gamma/\gamma_{\text{prod}}$ implies that maximizing the luminosity argues for accepting the muons into the muon linac at an energy somewhat lower than the energy which maximizes A_μ. The factor, C_μ, accounts for the possibility of cooling the muons; if no transverse cooling is used, $C_\mu = 1$ and $L_{\text{cool}} = 0$. At 100 GeV a reasonable value of β^* can be assumed to be 1 cm, although smaller values are possible, limited by the muon bunch length and by the design of the detector. Hence, the length of the muon bunch should be less than 1 cm. Such a short pulse is assured, if the length of the electron beam pulses are ≈ 0.5 cm.

5. Examples and parametric dependences

One now has a complete set of equations with which to maximize the luminosity of the muon collider as a function of the electron beam power incident on the production target and other system characteristics. As a first step in examining parametric dependences, we formulate a "realistic", baseline scenario that does not employ cooling of the muon beam.

The CLIC group at CERN [9] has developed a design concept for a high power positron production target to operate at 500 to 750 kW, more than an order of magnitude greater than presently operating designs. For the "realistic", baseline scenario assume that this target design can be realized at 0.5 MW. Using a 50 GeV electron beam with 20 nC per bunch and one bunch per macropulse, one can produce muon bunches of ≈ 0.1 nC at 29 GeV with an acceptance of ±3% in the capture section of the muon linac. The geometrical emittance of the muon beam at 29 GeV will be 5 π mm-mrad. If the average dipole field in the collider is 3 T, the revolution period will be 2 μs. Hence, the collision frequency will be ≈ 0.5 MHz. Then for β^* equal to 1 cm, the luminosity of the muon collider at 100 GeV will be ≈ 2×10^{26} cm^{-2}s^{-1}. This scenario, which we will use as a base case for parametric studies, is summarized as column 1 in Table 1 along with a more optimistic case without cooling (column 3).

The improvements to the "realistic" and "optimistic" cases that would obtain from damping the transverse emittance of the muons via ionization cooling are shown in columns 2 and 4 respectively. A far more optimistic scenario (column 5), which also requires several technological inventions including considerable cooling of the muon beam, is discussed in Section 6. In all the examples with beam cooling the ionization media are beryllium disks of thickness of $\beta_{\text{cool}}/2$. The beam is re-accelerated in rf-cavities with rf-cavities operating with an average accelerating gradient of 17 MeV/m.

The effect of the choice of the electron beam energy on the production efficiency can be seen in Fig. 6, which displays the maximum luminosity versus the electron beam energy for the realistic scenario. In this calculation the number of electron bunches is varied to keep the beam power on the muon production target fixed at 0.5 MW. The momentum acceptance is fixed at ±3%; however, the value of muon energy accepted and the angular spread of muons accepted is varied so as to maximize the luminos-

Table 1
Characteristics of a 100 GeV × 100 GeV muon collider using electro-production. The repetition rate in all cases is 500 Hz. For multiple rings, B_q, a_q, β_{cool} refer to the second ring. The quantities with daggers require technological inventions

	"Realistic" no cooling	"Realistic" with cooling	"Optimistic" no cooling	"Optimistic" with cooling	Needs inventions
Production					
E_e (GeV)	50	50	50	50	50
P_{beam} (MW)	0.5	0.5	2	2	5 †
N_e (nC)	20	20	30	30	50 †
E_{accept} (GeV)	29	21	29	22	25
$(\Delta p/p)_a$(%)	±3	±3	±4	±4	±8
N_μ (nC)	0.1	0.1	0.2	0.18	0.6
ϵ_a (π m rad)	1.95×10^{-3}	2.2×10^{-3}	1.95×10^{-3}	2.2×10^{-3}	3.0×10^{-3}
Cooler					
E_{cool} (GeV)	–	40	–	45	100 †
Number of rings	0	1	0	1	2
F_{cool}	–	0.5	–	0.5	0.6
$\langle B_d \rangle$ in (arc) (T)	–	4.5	–	4.5	4.5
V_{ring} (GeV/turn)	–	0.95	–	1.2	3.2
C_{ring} (m)	–	491	–	553	1840
(B_q (T), a_q (cm))	–	(4, 1.5)	–	(6, 1.2)	(8†, 0.5)
β_{cool} (cm)	–	1.3	–	1.5	0.4
ϵ_{med} (π m rad)	1.7×10^{-3}	5.7×10^{-5}	1.9×10^{-3}	7.4×10^{-5}	1.6×10^{-5}
C_v	1	38	1	28	136
Collider					
N_μ (nC)	0.1	0.068	0.2	0.14	0.35
N_{bunch}	1	1	2	2	2
B_{ave} (T)	3	3	4.5	4.5	6 †
$C_{collider}$ (m)	690	690	460	460	345
f_{coll} (MHz)	0.5	0.5	1.33	1.33	2
β (cm)	1	1	1	1	0.4
$(\Delta E/E)_{collider}$ (%)	±0.9	±0.8	±1.3	±1.0	±0.6
$\langle L \rangle$ (cm^{-2} s^{-1})	1.5×10^{26}	9.5×10^{26}	7.1×10^{26}	8.6×10^{27}	1.0×10^{30}

ity. As can be seen from Eqs. (2) and (22), the optimum acceptance energy will be a large fraction of the beam energy as the luminosity is quadratic in the conversion efficiency.

The scenarios employing production of muons at high initial energy (20–30 GeV) achieve relatively high luminosity at the expense of producing a muon beam with a relatively large ($\approx$ 1%) momentum spread at the interaction point. If a much lower spread, say ±0.1%, were required for physics reasons, then the accepted muon energy, E_{accept}, would have to be reduced to $\approx$ 5 GeV. The luminosity is still maximized by maximizing the electron beam energy. Making this change in E_{accept} to the "realistic" scenario reduces the luminosity to $\approx 6\times10^{24}$ $cm^{-2}s^{-1}$. As transverse cooling is accompanied by damping of the momentum spread, this consideration is not as severe in the scenarios with beam cooling.

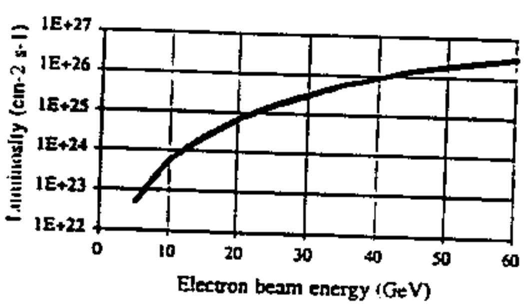

Fig. 6. The variation of collider luminosity with energy of the electron beam at the production target in the "realistic" scenario. The beam power is fixed at 0.5 MW.

The optimum energy for accepting the muons in the absence of cooling is 29 GeV. If instead we employ an ionization cooler, the optimum acceptance energy would be reduced to 21 GeV; a curve of the luminosity versus muon acceptance energy for the "realistic scenario" is given in Fig. 7. In this calculation cooling energy has been optimized but limited to $\leq$ 40 GeV.

Somewhat surprisingly, the higher the energy at which the muons are cooled, the higher the final luminosity. The reason is that the adiabatic damping of the energy spread permits a much smaller value of β_c. If an initial stage of

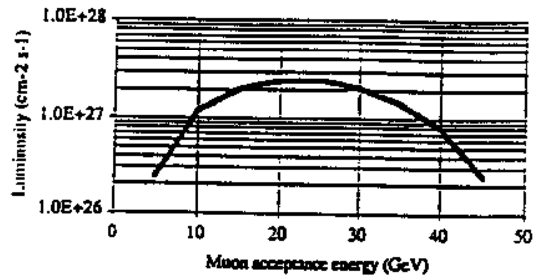

Fig. 7. Luminosity as a function of muon acceptance energy for the "realistic" scenario with ionization cooling.

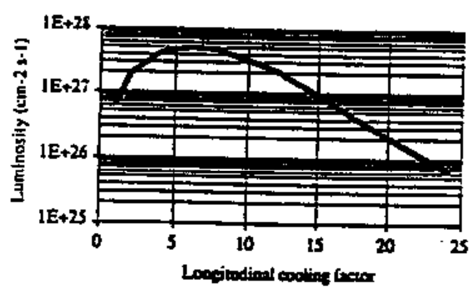

Fig. 8. Luminosity variation with longitudinal cooling for the "realistic" scenario.

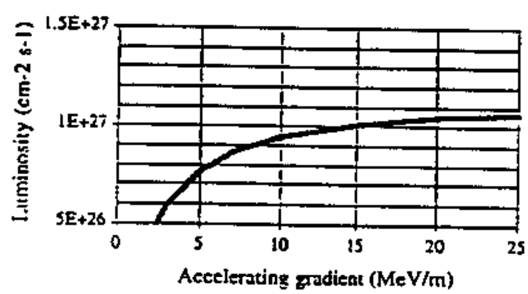

Fig. 9. Luminosity versus accelerating field in the cooling ring for the "realistic" scenario of Table 1.

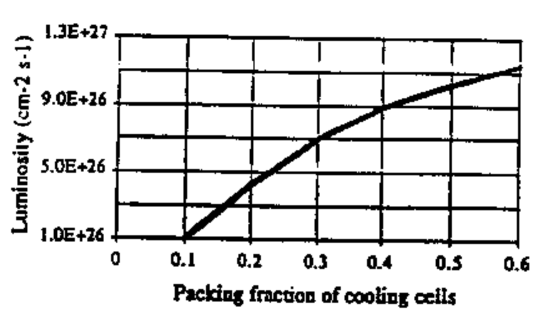

Fig. 10. Luminosity versus packing fraction of ionization cells and re-acceleration cavities in the cooling ring for the "realistic" scenario of Table 1.

ionization cooling were employed to reduce the energy spread, the optimum energy at which transverse cooling is performed could shift to a lower value. In the "realistic" example, the traverse cooling by a factor of 48 at 40 GeV requires an energy exchange of only $3.8E_c$. From Eq. (17) one expects a slight improvement in C_μ from the damping of the energy spread in the zero-dispersion cells. Adding ionization cells in the dispersive sections of the ring as suggested in Ref. [6] could improve the luminosity significantly.

Fig. 8 illustrates the variation in luminosity for the "realistic" scenario with longitudinal cooling accompanying the transverse emittance damping. In this calculation cooling is done at the muon acceptance energy, 21 GeV so that no additional adiabatic reduction in energy spread is included. The field strength and aperture of the cooling channel optics is kept fixed. The decrease in luminosity as the cooling factor increases beyond six comes from the decay of the muon population as the transverse cooling path length increases to allow the beam to reach the equilibrium emittance.

As the muons must remain in the cooling lattice for hundreds of microseconds, it may not be possible to maintain an accelerating gradient of 17 MeV/m as assumed in the examples of Table 1. The consequence of reducing the gradient to allow for a lower power accelerating system in the cooling ring is displayed in Fig. 9. The degradation of the luminosity becomes especially large as the gradient falls below 10 MeV/m. As the number of muons in the ring is small, the beam loading in the cooling ring will be very small. One might consider the use of superconducting rf-cavities to keep rf-power requirements relatively small. Whether the superconducting cavities can function in the presence of radiation from the muon decay is uncertain.

A second characteristic of the ionization cooling lattice that can have a strong effect on the final luminosity of the collider is the packing fraction, F_c, of the ionization cells plus the rf-cavities in the cooling ring. Fig. 10 illustrates the variation of luminosity with F_c for the "realistic" case with cooling.

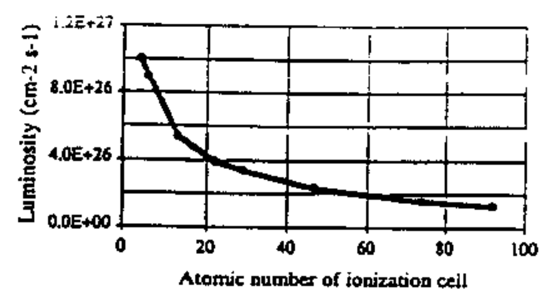

Fig. 11. Variation of luminosity with the choice of ionizing medium.

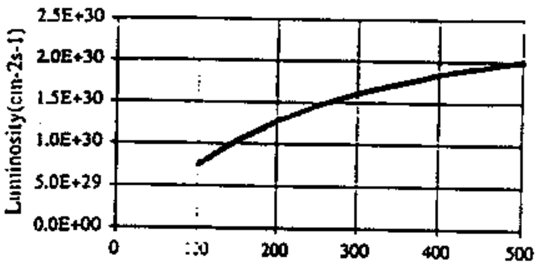

Fig. 12. Variation of luminosity with energy for a 250 GeV×250 GeV muon collider with $\beta^* = 0.3$ cm.

Applying Eq. (7) through Eq. (17) to calculate the characteristics of a cooling system, we find that the luminosity varies with choice of the ionizing medium as shown in Fig. 11. Although the product X_R dE/dx is independent of density, the luminosity is sensitive to the density of the ionizing medium as the energy lost per cell depends on the density and thickness of the medium. For each of the points in Fig. 11 the appropriate density has been used. From this examination we confirm that the preferred ionization media are beryllium disks.

If one were to design the muon collider with a broader energy reach, for example from 100 to 500 GeV center of mass energy, one would hope to realize a higher luminosity at the higher energies as the geometrical emittance is reduced by adiabatic damping. The scaling of the luminosity, as shown in Fig. 12, is slower than linear. The calculation of Fig. 12 is based on the "Needs invention" scenario of Table 1, with β^* reduced to 0.3 cm.

In this scenario the momentum spread of the beams is largest at the lowest energy. Unfortunately, the width of a Standard Model Higgs is expected to be a rapidly increasing function of the Higgs mass with a value of 1 GeV/c^2 for $m_H = 100$ GeV/c^2. If the momentum spread were reduced at the lower energies to allow for a fine scan of the range from 100 to 200 GeV, the luminosity would fall off much more precipitously.

6. Prospects and conclusions

To obtain a muon collider with a luminosity of 10^{30} cm^{-2}s^{-1}, as desired for studies of the Higgs, one must adopt an extremely optimistic scenario (column 5 of Table 1) that includes several technological innovations (indicated by a dagger). Perhaps the easiest of these advances may be the development of very strong, precision dipoles that would enable one to design a relatively small storage ring collider with a dipole field of 6 T averaged over the entire ring.

Continuing advances in the technology of electron guns with photocathodes suggest that one may be able to obtain 50 nC bunches of electrons for injection in a S-band structure. Accelerating multiple bunches of such high charge in a S-band structure also presents difficulties. For a 50 nC, 15 ps bunch, the single-bunch beam loading in a SLAC structure operating at 20 MeV/m would be ≈ 20%. As the bunches in the macropulse are separated by hundreds of meters, multi-bunch beam breakup is not a problem. However, a head-to-tail momentum variation of 4% will be required for BNS damping of the single bunch, transverse, head-to-tail instability. Once this systematic variation is removed at the end of the electron linac, one would be left with a ±1.5% spread that must be handled by the focusing optics at the production target.

Extending the conceptual design of the CERN production target to a reliable, 5 MW design is likely to be very difficult. Of particular difficulty will be finding suitable accelerator components that can withstand the extremely high radiation environment near the target. Note that the highest power, production target in operation is the 33 kW positron production target at SLAC.

It is likely that the greatest challenge to the designer will be to find an efficient scheme for cooling the muon beam at a high initial energy. In scenarios that include beam cooling in a storage ring the momentum bite must be chosen to be consistent with the acceptance of the cooling lattice. As it should be possible to design a lattice with an acceptance of ±2%, cooling the muons at very high energy allows accepting a large momentum bite at the production target.

An idea of the scope of the project can be had by observing that in the realistic case the collider ring has a circumference of ≈ 690 m while the cooler rings (one each for the μ^+ and μ^-) have circumferences of ≈ 790 m. Operating with a gradient of 17 MeV/m, the electron linac would be 3 km long while the 20 GeV muon linac would have a length of 1.32 km. A clever design may be possible in which these same linacs could be used to accelerate the muons from the cooler ring up to the full 100 GeV per beam of the collider. In this case the major cost of the project would be the 70 GeV of S-band linac. The major complexity and technological risk is in the lattices of the cooling rings which use very high field, superconducting quadrupoles and dipoles.

In conclusion, one sees that even with optimistic assumptions, it is difficult to envision a high energy $\mu^+\mu^-$ collider which employs electro-production of muons functioning with a luminosity $> 10^{27}$ cm^{-2}s^{-1}. While the possibility of an electron-beam-driven muon collider with a luminosity ≈ 10^{30} cm^{-2}s^{-1} cannot be ruled out, it would require major advances in several of the primary constituent technologies. The areas for innovations include superconducting dipoles and quadrupoles, multi-kilo-ampere electron beam sources, and multi-megawatt muon production targets. Most critically, efficient means of both transverse and longitudinal of cooling the muon beams at high energy must be found and demonstrated, if suitably high luminosity is to be achieved.

References

[1] D.B. Cline, Opening presentation at the Workshop on High Energy Muon Colliders, Napa, California, December, 1992.
[2] W.R. Nelson, K.R. Case and G.K. Svensson, Nucl. Instr. and Meth. 120 (1974) 413.
[3] W.R. Nelson, Nucl. Instr. and Meth. 66 (1968) 293.
[4] W.A. Barletta, and A.M. Sessler, Radiation from fine self-focussed beams at high energy, in: High Gain, High Power Free Electron Lasers; eds. L. Bonifacio, L. De Salvo-Souza and C. Pellegrini (Elsevier, 1989).
[5] A. Belkacem et al., Experimental study of pair creation and radiation in Ge crystals at ultra-relativistic energies, in: Relativistic Channeling., eds. Carrigan and Ellison, Nato ASI Series B, Vol. 165 (Plenum).
[6] D. Neuffer, Part. Accel. 14 (1983) 75.
[7] Particle Data Group, Phys. Rev. D45 Part 2 (1992).
[8] K. Brown, 'Private communication as reported in R. Palmer, Interdependence of Parameters for TeV Linear Colliders, Proc. Workshop on New Developments in Particle Accelerator Techniques, Orsay, (June 1987).
[9] P. Sievers, Proc. Workshop on Heavy Quark Factories, Courmayeur, Italy, December, 1987.

Nuclear Instruments and Methods in Physics Research A 350 (1994) 45–52
North-Holland

NUCLEAR
INSTRUMENTS
& METHODS
IN PHYSICS
RESEARCH
Section A

A muon collider scenario based on stochastic cooling ☆

Alessandro G. Ruggiero

Department of Advanced Technologies, Brookhaven National Laboratory, Upton, Long Island, New York 11973, USA

Received 8 April 1994

The most severe limitation to the muon production for a large-energy muon collider is the short time allowed for cooling the beam to dimensions small enough to provide reasonably high luminosity. The limitation is caused by the short lifetime of the particles that, for instance, at the energy of 100 GeV is of only 2.2 ms. Moreover, it appears to be desirable to accelerate the beam quickly, with very short bunches of about a millimeter so it can be made immediately available for the final collision.

This paper describes the requirements of single-pass, fast stochastic cooling for very short bunches. Bandwidth, amplifier gain and Schottky power do not seem to be of major concern. Problems do arise with the ultimate low emittance that can be achieved, the value of which is seriously affected by the front-end thermal noise.

Since mixing within the beam bunches is completely absent, methods are required for the regeneration of the beam signal with external and powerful magnetic lenses. The feasibility of these methods are crucial for the development of the muon collider. These methods will be studied in a subsequent report.

1. Introduction

In the quest for the Higgs bosons, a muon collider may be perceived as the experimental device more affordable and more feasible than electron–positron or very large hadron colliders [1–3]. Muons have a mass ten times lighter than protons and are therefore easier to be steered on circular trajectories. On the other hand their mass is a hundred times greater than electrons and their motion is considerably less affected by the synchrotron radiation. Muons are elementary lepton particles, with no internal structure. Like the electrons, they have obvious advantages over the hadron counterpart when they are used as the main projectiles for the production of the Higgs bosons. Moreover, because of their larger mass, they are also better suited than the electrons themselves, due to a considerably larger propagator constant. Unfortunately, muons do not exist in nature and they have to be produced with the only technique we know these days: impinging an intense beam of protons or electrons on a target. Like in the case of production of antiprotons, in order to make the beam of some use for the subsequent collisions, muons also have to be collected and cooled to a sufficiently high intensity and small dimensions before they can be accelerated and injected in the collider proper. To make the situation more complicated there is also the fact that muons are intrinsically unstable particles with a very short lifetime. Accumulation, acceleration and cooling are then to be executed extremely fast if one requires that a large fraction of the particle beam survives to the collision point.

This paper deals with the requirement of betatron stochastic cooling which is to be very effective and fast. The situation being described is altogether different from the usual encountered with coasting beams [8,9]. Now the beam is tightly bunched at a very large frequency. The bunches are very narrow, having a length which is considerably smaller than the wavelength of the bandwidth of available electronic amplifiers. Thus a different method is to be developed based on the correction of the stochastic signal for all particles at the same time in one single-step. The fundamental limitation remains of the ultimate value of the final emittance that can be achieved. The limitation is caused by the thermal noise at the front-end of the amplifier.

We begin by reviewing a possible scenario of a muon collider in Section 2. This scenario assumes that stochastic cooling is done at maximum energy. The performance of the collider luminosity is evaluated in Section 3, where special emphasis is put on the effects of the beam lifetime and on the betatron emittance reduction. The requirements for the stochastic cooling proper are exposed in Section 4, where we underline the particular situation we are facing of very short beam bunches. The analysis of the cooling device itself follows in Section 5. The goal is the determination of an equation which gives the evolution of the betatron emittance with time. This is described essentially by two parameters: the cooling rate and the diffusion rate due to the thermal noise which is by far more important

☆ Work performed under the auspices of the U.S. Department of Energy.

SSDI 0168-9002(94)00522-9

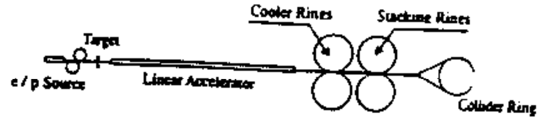

Fig. 1. A conceptual layout of the muon collider.

than any other diffusion process, for instance beam Schottky power. The derivation of the equation for the evolution of the betatron emittance is given in Section 6. A discussion leading to optimization considerations of the cooling process is presented in Section 7. Finally, an application to the muon collider is worked out in Section 8, where the dependence of the performance with a dynamical gain less than optimal and with the bunching frequency is also investigated. It is found that a luminosity of about 10^{24} cm^{-2} s^{-1} can be achieved at the very most. Conclusions are given in the last Section 9.

2. The muon collider

We expose below a possible scenario of a muon collider. A layout of the scheme is shown in Fig. 1. An intense source of either protons or electrons, at the energy of few tens of GeV with an average current of few hundreds of μA, is provided with a conventional fast cycling accelerator [4–8]. In the case of electrons, these are bunched at a large frequency, for instance 3 GHz and are accelerated in a large-gradient linear accelerator. In the case of protons, after acceleration, the beam is debunched in a stretcher ring and then rebunched also at the frequency of 3 GHz. In either case, the primary beam is made to impinge on a sequence of targets for the production of muon pairs via decay of π mesons. The secondary beam, produced at an energy of about 1 GeV, will be collected with a large production angle yielding a normalized total emittance of about 100 π mm mrad, and with a large momentum bite of few percent. An average intensity of about 50 nA per each component of the pair production is expected, corresponding to a yield of approximately 2×10^{-4} μ-pairs per primary particle. The μ-beams are also tightly bunched at the frequency of 3 GHz so that there are only few particles per bunch (around 100). The lower number of particles per bunch is, as we shall see, a requirement for fast stochastic cooling.

Both types of beam, μ^+ and μ^-, after a preliminary bunch rotation to reduce the momentum spread, are accelerated in a large-gradient linear accelerating structure, operating also at 3 GHz, to the final energy which is in the range of 100 to 1000 GeV. At the end of the acceleration, each beam is transferred to a storage ring where fast stochastic cooling is done to reduce the betatron emittance to the final value. Each beam is then taken to a stacking ring of about the same size, where several cooled beam pulses are stacked side-wise in the momentum phase space. The stacking procedure is carefully done to avoid lengthening of the bunches and increasing of the betatron emittance. With this operation the number of particles per bunch will increase by the number of pulses being stacked. This is required to boost the magnitude of the luminosity of the collider [7]. At the end of stacking, both beams are extracted from their respective stacking ring and transferred to the collider ring proper where they are made to collide and exploited for experimentation. It is to be noticed that all storage rings involved (five, as shown in Fig. 1) operate at the same energy; thus it is expected that they have also about the same size.

Because of the relatively short lifetime of the muon particles, it is obvious that all the operations which have been described above are to be executed very fast.

3. The luminosity performance

What follows is a discussion of the luminosity performance of the muon collider. The average luminosity is given by the following expression

$$L = MN_0^2 f_{\mathrm{bunc}} F\gamma / 4\pi\epsilon_n \beta^*, \tag{1}$$

where $N_0 \sim 100$ is the initial number of particles per bunch at the moment of production, $f_{\mathrm{bunc}} \sim 3$ GHz is the beam bunching frequency during acceleration and stochastic cooling, γ is the energy relativistic factor, $\epsilon_n \sim 25$ π mm mrad is the initial rms normalized emittance, at the moment of production, and β^* is the focussing amplitude parameter at the interaction point, which for a multiple pass in a collider ring can be as low as 10 cm and for the single-pass mode, where the requirements on the lattice focussing can be relaxed, it is about 1 cm. For an efficient mode of operation, it is important that the bunch length during collision is sufficiently small when compared to β^*. This is obtained with the large bunching frequency. M is the number of beam pulses which are stacked in the momentum phase space of the stacking ring. It is to be noticed that the current $I_\mu = N_0 e f_{\mathrm{bunc}} \sim 50$ nA is a constant equal to the average current of each muon beam at the moment of production. The luminosity expression above shows clearly the advantage of increasing the number of particles per bunch from N_0 to MN_0 with momentum stacking [7]. Finally F is a form factor which includes the losses of particles due to the short lifetime and the emittance reduction due to the stochastic cooling.

We can write the form factor F as the product of many other factors:

$$F = F_{\mathrm{acc}} F_{\mathrm{sto}} F_{\mathrm{stac}}^2 F_{\mathrm{col}}, \tag{2}$$

where F_{acc} is the square of the beam survival fraction after acceleration, F_{sto} reflects the effects on the luminosity of the reduction of the betatron emittance due to stochastic cooling and of the square of the fraction of beam survival after cooling, F_{stac} is the beam survival fraction after

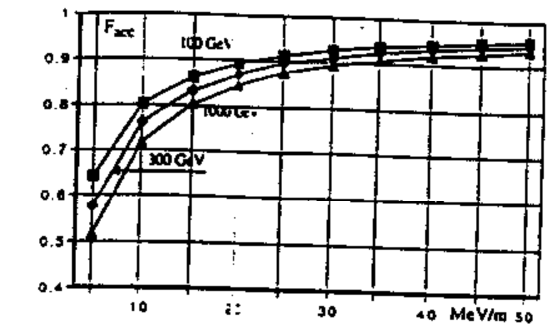

Fig. 2. The acceleration survival factor F_{acc} vs. acceleration gradient.

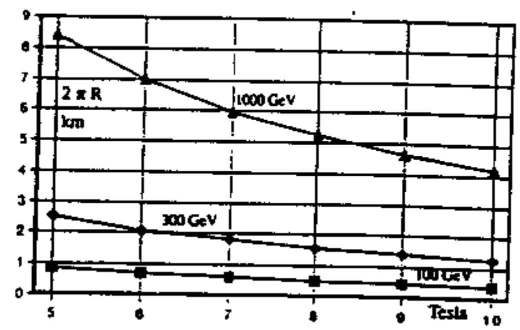

Fig. 4. Circumference of storage ring vs. bending field B for $\eta = 0.5$.

momentum stacking, and F_{col} represents the beam losses during collision.

The overall survival ratio after acceleration can be easily calculated by integrating the instantaneous beam survival over the acceleration cycle, combined with the fact that the lifetime increases linearly with γ. Taking into account the contribution of two beams to the luminosity, we have

$$F_{acc} = (E_{nit}/E_{final})^{2E_0/cG\tau_0}, \tag{3}$$

where $E_{init} \sim 1$ GeV is the beam kinetic energy at production, E_{final} is the final energy in the collider, $E_0 = 106$ MeV is the rest energy, $\tau_0 = 2.2$ μs the lifetime of the muon at rest and G is the accelerating gradient in the linear accelerator. The behavior of F_{acc} versus the accelerating gradient is shown in Fig. 2 for various final energies. It is seen that losses are reduced with a larger accelerating gradient and a lower beam energy.

The second factor F_{sto} represents the combined effect of the beam losses during the period of time T_{sto} the beam spends to be cooled and the reduction of the betatron emittance with stochastic cooling

$$F_{sto} = (\epsilon_0/\epsilon_x)\exp(-2T_{sto}/\gamma_{final}\tau_0), \tag{4}$$

where (ϵ_0/ϵ_x) is the ratio of the initial to the final betatron emittance.

The factor F_{stac} represents the survival of the beam integrated over the M pulses being stacked in the stacking ring over a period of time T_{stac}. By denoting with f_0 the circulating frequency, we have $T_{stac} = M/f_0$ and

$$F_{stac} = [1 - \exp(-T_{stac}/\gamma_{final}\tau_0)]/(T_{stac}/\gamma_{final}\tau_0). \tag{5}$$

This quantity is plotted in Fig. 3 versus the parameter $\zeta = T_{stac}/\gamma_{final}\,\tau_0$.

The circumference $2\pi R$ of any of the storage rings can be estimated from the beam energy γ_{final}. Allowing a bending-magnet packing-factor η, denoting with B the banding field in T, and expressing the circumference in meter gives

$$2\pi R = 2.22\gamma_{final}/\eta B \tag{6}$$

from which we derive the revolution frequency

$$f_0 = (135\ \text{MHz})\eta B/\gamma_{final}. \tag{7}$$

Thus is seen that

$$\zeta = 0.0034\ M/\eta B, \tag{8}$$

which does not depend on the beam energy. If we take $B = 6$ T and $\eta = 0.5$ then $\zeta \approx 0.0011M$. To avoid excessive luminosity losses, it is seen from Fig. 3 that at most we can allow the momentum stacking of $M \sim 900$ pulses. The quantity ζF_{stac}^2 is also plotted in Fig. 3 which shows that at most $MF_{stac}^2 \sim 350\eta B$.

For completeness we display the plots of the circumference $2\pi R$ and of the revolution frequency f_0, respectively

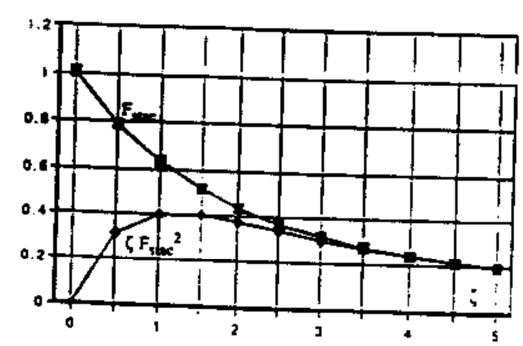

Fig. 3. The survival factors F_{stac} and F_{stac}^2 vs. the parameter.

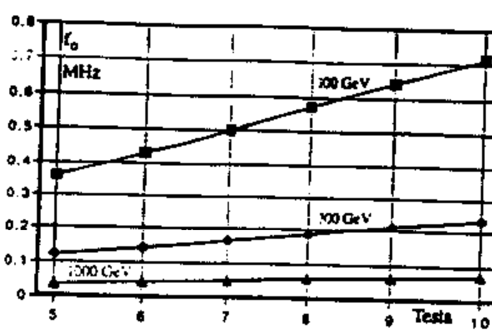

Fig. 5. The revolution frequency vs. bending field B ($\eta = 0.5$).

in Figs. 4 and 5, versus the bending field B and for various beam energy values.

Finally, F_{col} represents the loss of average luminosity due to the particle losses during collision which we assume takes a period of time T_{col}. This factor has an expression similar to the one for F_{stac} given by Eq. (5), except that T_{stac} is replaced by $2T_{col}$. It is seen that the best case is given by the single pass mode where after one interaction both beam bunches are immediately disposed. This mode is also more favorable because it does not require a complete collider ring and the final focus may correspond to a lower value of β.

4. Requirements on stochastic cooling

We can express the actual average luminosity in terms of the ideal value L_0 without stochastic cooling, without momentum stacking and for an infinitely long muon lifetime

$$L = MFL_0, \tag{9}$$

where

$$L_0 = N_0^2 f_{bunc} \gamma / 4\pi\epsilon_n \beta^*. \tag{10}$$

With the values of the parameters given in the previous two sections and for the single-pass mode, which is the most favorable, we have $L_0 = 1.0 \times 10^{15}\gamma$ cm^{-2} s^{-1}. Correspondingly, to achieve a luminosity of about $1.0 \times 10^{27}\gamma$ cm^{-2} s^{-1}, which is required for the high energy physics experimental exploitation of the collider, we need that the overall enhancement factor $MF \sim 1 \times 10^{12}$. Since at most $MF_{stac}^2 \sim 1000$, even assuming $F_{acc}F_{col} \sim 1$, we need $F_{sto} \sim 1 \times 10^9$, which is a very large requirement for the betatron stochastic cooling. The requirement is independent of the beam energy; a normalized emittance of $25 \times 10^{-9}\pi$ mm mrad is required at all energies. Assuming a cooling period T_{sto} short with respect to the muon lifetime $\gamma_{final}\ \tau_0$, one requires a reduction of the betatron emittance by nine orders of magnitude (!). Thus the fundamental question concerns the ultimate emittance that can be realistically achieved at the end of cooling.

There are major differences between stochastic cooling for the case of bunched beams we are investigating here and the usual approach for coasting beams encountered, for instance, during production and accumulation of antiproton beams [9,10]. The muon bunches have a length considerably smaller than the shortest wavelength in the frequency bandwidth of the system. It is indeed a good approximation to assume that the beam bunches have no longitudinal extension, and that all the particles are distributed on a disk with a center slightly displaced from the axis of the pickups. The beam current signal is therefore highly organized and coherent. The transverse beam position, on the other hand, has a very stochastic behavior. Because of the low number of particles, there is a random fluctuation of the beam centroid that can be related statistically to the overall transverse beam size. For the same reasons, the internal motion can be completely ignored and no mixing occurs between the detection of the beam signal at the pickups and the application of the deflection at the kickers. Moreover, the conventional analysis in the frequency domain [9,10] would be hardly applicable, since the Schottky bands remain well separated from each other. Actually from the current point of view, it is improper to refer to the beam Schottky signal. As pointed out already, the stochastic behavior appears only in the transverse displacement of the beam centroid.

The overall system, between pickups and kickers, including the power amplifier, has a bandwidth W large enough to detect and correct displacement of individual bunches. If we assume a bandwidth W extending over an octave, where the frequency at the upper end is twice than the frequency at the lower end, an optimum is given by choosing the bandwidth equal to an integer m times 2/3 of the bunching frequency. For instance, if the beam is bunched at 3 GHz, the bandwidth could be 2 GHz ($m = 1$) extending from 2 to 4 GHz. A larger bandwidth, for instance from 4 to 8 GHz ($m = 2$) or from 6 to 12 GHz ($m = 3$), is of course also possible. In this mode of operation, it is possible to process the pickup signals to allow complete rejection of a beam bunch signal on another bunch.

The lack of mixing causes a serious limitation on the effectiveness of stochastic cooling. Once the initial beam displacement has been corrected, there is no more signal from the beam that can be used. Thus everything is done in a single step with a relatively small reduction of the beam size. Between steps, the signal from the beam has to be regenerated, for instance by rearranging the particle mutual position with the aid of powerful magnetic lenses. We shall assume below that this is indeed the case. We shall investigate in a separate report the amount of particle rearrangement required and how this can be realized in practice.

Because of the low number of particles per bunch, after amplification, the beam Schottky power is not expected to be excessive. In order to obtain a very fast cooling, one requires to optimize the overall gain to correct the instantaneous beam displacement immediately in one single step. This may require a very large electronic gain. A more serious problem is associated with the thermal noise at the front-end of the amplifier, which will set a limitation of the final beam transverse dimension. The design of the cooling device is to be optimized to reduce this limitation.

5. Analysis of the cooling device

Consider a very narrow bunch made of N particles all with the same electric charge. The bunch is periodically traversing a beam position pickup made of two parallel striplines each of length l and separated by a distance d. The striplines are shorted at one end and terminated at the upstream end to their characteristic impedance R_p. During the occurrence of a traversal assumed at the instant $t=0$, the bunch current can be represented as a pulse of zero duration, proportional to the average displacement $\bar{x}$, that is

$$I_p(t) = Ne(\bar{x}/d)\delta(t). \quad (10)$$

This current leaves at the upstream end of one pickup a voltage signal given by

$$V_p(t) = Ne(\bar{x}/2d)R_p[\delta(t) - \delta(t-2l/c)], \quad (11)$$

that is a voltage pulse occurring simultaneously to the current pulse, followed by another at the delay of $2l/c$. Since the beam bunch duration is considerably shorter than the delay between the two voltage pulses, only the first pulse is relevant to our analysis and we shall ignore the second one which we assume can be disposed properly without disruptions to the subsequent bunches. The voltage signal is then filtered by the bandwidth of the system, mostly caused by the power amplifier, and properly amplified by the linear gain A. The resulting voltage is an oscillating and decaying signal of which only the front-end is of relevance here since it constitutes the part that is to be applied in phase with the bunch at the location of the kickers. The amount of the properly correspondent voltage is simply

$$V_A = Ne(\bar{x}/2d)R_p AW. \quad (12)$$

In the case there is only one pickup and one kicker, this is also the voltage that would appear across the kicker assuming an ideal impedance matching all along the transfer of the signal. To include the case of more pickups and kickers, we modify Eq. (12) as follows:

$$V_k = Ne(\bar{x}/2d)AW\sqrt{R_p R_k n_p/n_k}, \quad (13)$$

where $n_{p(k)}$ is the number of pickups (kickers) and R_k is the characteristic impedance of kickers. This expression gives the voltage across one single kicker. We are assuming here that both the pickups and the kickers are closely packed and that they extend over a length of the storage ring where the beam position $\bar{x}$ and the lattice functions do not vary appreciably. Moreover kickers have exactly the same geometrical configuration and size of the pickups.

Assuming small deflection angles and full correlation among the kicks, the total deflection angle each particle will be subject to at the traversal of the kickers is

$$\theta_S = eV_k l n_k/\beta^2 Ed, \quad (14)$$

where E is the particle total energy and βc is the velocity. Statistically, over many revolutions, the following relation holds between the average beam displacement $\bar{x}$ and the rms beam size σ:

$$\sigma^2 = N\bar{x}^2. \quad (15)$$

We prefer writing the expression for the deflection θ_S as follows:

$$\delta_S = g_0\sigma/d, \quad (16)$$

where

$$g_0 = \frac{lA}{2\beta^2 Ed}\sqrt{Ne^4W^2 n_k n_p R_k R_p}. \quad (17)$$

Another signal is induced to the kickers which is random and independent of the particle position. The finite temperature of the terminating resistors of the loop and of the preamplifiers creates at the input to the preamplifier a signal of power

$$P_T = k_B(T_A + T_R)W, \quad (18)$$

where $k_B = 8.6171 \times 10^{-5}$ eV/K is the Boltzmann constant, T_A is the equivalent temperature of the amplifier and T_R is the temperature of the resistor. Proceeding in the same way as for the beam Schottky signal, we calculate the total deflection angle due to the thermal noise:

$$\theta_T = \frac{elA}{\beta^2 Ed}\sqrt{n_k R_k P_T}. \quad (19)$$

To assess the effectiveness of stochastic cooling a useful parameter is the ratio of Schottky to the thermal power

$$S = \theta_S^2/\theta_T^2 = Ne^2W^2 n_p R_p \sigma^2/4d^2 P_T. \quad (20)$$

It is seen that this ratio increases linearly with the bandwidth. This result is different from the one it was derived for coasting beams [10], in which case the ratio is independent of the bandwidth. The difference is due to the fact that now the current distribution of the beam bunch is highly organized and does not exhibit a stochastic behavior.

6. The equation for the evolution of the beam emittance

Both Schottky and thermal power kicks apply simultaneously when a particle is crossing the location of the kickers. Let us calculate the effect of both kicks to the beam emittance which we can define as follows

$$\epsilon = \sum_i \left(\gamma x_i^2 + 2\alpha x_i x_i' + \beta x_i'^2\right)/N. \quad (21)$$

Here α, β and γ are the lattice Twiss parameters.

At the kickers each particle receives the same kick, that is $x_i' \to x_i' + \theta_S + \theta_T$. The corresponding emittance change is

$$\Delta\epsilon = 2(\alpha\bar{x} + \beta\bar{x}')_k(\theta_S + \theta_T) + \beta_k(\theta_S + \theta_T)^2, \quad (22)$$

where the subscript k(p) denotes that the corresponding quantity is evaluated at location of the kickers (pickups). $\bar{x}$ and $\bar{x}'$ are the average values of the particle positions and angle. We are interested in the expectation value of $\Delta\epsilon$

over many revolutions. In this case, since there is no correlation between particle position and thermal noise, the previous equation reduces to

$$\langle \Delta\epsilon \rangle = 2(\alpha\bar{x} + \beta\bar{x}')_k \theta_S + \beta_k \theta_S^2 + \beta_k \theta_T^2. \tag{23}$$

The betatron emittance is also defined as

$$\epsilon = (\sigma^2/\beta)_{p,k}. \tag{24}$$

At the same time it can be proven that

$$\langle (\alpha\bar{x} + \beta\bar{x}')_k \bar{x}_p \rangle = -\langle \bar{x}_p^2 \rangle \sqrt{(\beta_k/\beta_p)} \sin\psi_{pk}, \tag{25}$$

where ψ_{pk} is the betatron phase advance between pickup and kicker. Manipulating some of the previous equations gives

$$\langle \Delta\epsilon \rangle = -(2g \sin\psi_{pk} - Ng^2)\epsilon + \beta_k \theta_T^2 \tag{26}$$

with the dynamical gain

$$g = g_0 \sqrt{\beta_k \beta_p / Nd^2}. \tag{27}$$

Finally, the beam emittance evolution is described by the following equation

$$\mathrm{d}\epsilon/\mathrm{d}t = -\lambda\epsilon + D, \tag{28}$$

where the cooling rate

$$\lambda = n_S f_0 (2g \sin\psi_{pk} - Ng^2) \tag{29}$$

and the diffusion coefficient

$$D = n_S f_0 \beta_k \theta_T^2. \tag{30}$$

with f_0 the revolution frequency. In deriving these equations we have assumed a total of n_s identical cooling systems in the storage ring. In the following we shall assume that the distance between pickups and kickers is adjusted so that $\sin\psi_{pk} = 1$.

7. Optimization of the cooling performance

An optimum cooling rate is obtained by setting $g = 1/N$ and is

$$\lambda_{opt} = n_S f_0 / N. \tag{31}$$

which corresponds to correcting the instantaneous beam bunch displacement in one single step. At the same time we can also derive the required amplifier gain

$$A = \frac{2\beta^2 E d^2 / lN}{\sqrt{e^4 W^2 n_k n_p R_k R_p \beta_k \beta_p}}, \tag{32}$$

which decreases linearly with the bandwidth and the number of particles in the bunch.

The equilibrium value of the emittance for this case is

$$\epsilon_x = D/\lambda_{opt} = N\beta_k \theta_T^2. \tag{33}$$

By combining some of the equations we obtain for the equilibrium emittance

$$\epsilon_x = \frac{4d^2 P_T}{N e^2 W^2 n_p \beta_p R_p}, \tag{34}$$

which, similarly to the amplifier gain A, also exhibits the same dependence with bandwidth W and number of particles N. It can be seen that the equilibrium emittance corresponds to the situation where the ratio of the Schottky to thermal power $S = 1$.

It is to be noticed that an optimum bandwidth is related to the bunching frequency f_{bunc} by the relation

$$W = 2mf_{bunc}/3, \tag{35}$$

where m is a positive integer. Also the quantity

$$I_\mu = Nef_{bunc} \tag{36}$$

is the average muon current essentially equal to the one produced at the target. Thus when these relationships are taken into account we have for the amplifier gain

$$A = \frac{3\beta^2 E d^2 / l I_\mu m}{\sqrt{e^2 n_k n_p R_k R_p \beta_k \beta_p}} \tag{37}$$

and for the equilibrium emittance

$$\epsilon_x = \frac{6d^2 P_T / W}{m n_p \beta_p e I_\mu R_p}. \tag{38}$$

Both of these expressions show the same dependence on the bandwidth factor m and on the average beam current. Noticing that the thermal power P_T is proportional to W, it is seen that both A and ϵ_x do not depend explicitly on how the beam is bunched. On the other hand the cooling rate λ depends very strongly with the number N of particles per bunch.

8. An application of the optimal system

As an application of these expressions we take the following values:

$d = 1$ cm,

$\beta_p = \beta_p = 200$ m,

$n_p = n_k = 1024$,

$R_p = R_k = 100\ \Omega$.

We chose $m = 3$, that is a bunching frequency $f_{bunc} = 3$ GHz, corresponding to a bandwidth $W = 6$ GHz ranging between 6 and 12 GHz. The length of the pickups is adjusted to match the bandwidth according to $l = 6$ cm/W (GHz) so that for $m = 3$ it is $l = 1$ cm. We assume also a bending field $B = 6$ T and a packing factor in the storage ring $\eta = 0.5$. Finally we set the temperature of the amplifier and resistor $T_A = T_R = 1$ K which is very likely an unrealistic value, for which moreover we cannot really

Table 1
Stochastic cooling performance

Beam energy, GeV	100	300	1000
$2\pi R$, m	700	2100	7000
n_s	8	24	80
$1/\lambda$, ms	0.030	0.030	0.030
A	1×10^9	3×10^9	1×10^{10}
$\epsilon_n = \gamma\epsilon_\infty$, π mm mrad	32	96	320
L_0, cm^{-2} s^{-1}	1×10^{18}	3×10^{18}	1×10^{19}
MF	1000	300	100
L, cm^{-2} s^{-1}	1×10^{21}	1×10^{21}	1×10^{21}
Ng_{max}	0.0068	0.0023	0.0007
f_{max}	40	120	400
L_{max}, cm^{-2} s^{-1}	4×10^{22}	1.2×10^{23}	4×10^{23}

foresee the behavior of the thermal noise at the front end. The summary of the results of our calculations are shown in Table 1, where we have also assumed the optimum gain $g = 1/N$.

To be observed is the increase of the circumference of the storage ring with the beam energy, and that we let the number n_s of cooling systems vary proportionally. As a consequence, the cooling time $1/\lambda$ is constant with energy, whereas the amplifier gain A and the equilibrium emittance ϵ_∞ increase linearly with energy. As one can see, even at the very low temperature of 1 K, thermal noise dominates over the beam signal, and the equilibrium emittance is just about comparable to the initial beam emittance at the energy of 100 GeV. For larger energies there is actually stochastic heating accompanied by an increase of the beam emittance.

Since for the optimum gain the cooling time is 0.03 ms which is considerably shorter than the beam lifetime, it is reasonable to lower the amplifier gain. The results are shown in Figs. 6–8 where the cooling rate, the amplifier gain and the equilibrium emittance are plotted versus the dynamical gain $g/g_{opt} = Ng$.

It is seen that, as the gain g is lowered, both the amplifier gain and the equilibrium emittance reduce also, but on the other hand, unfortunately, the cooling time increases. If the increase is too large then the particle losses would also be too large.

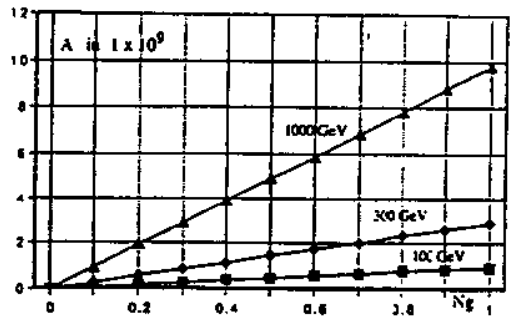

Fig. 7. Amplifier gain vs. dynamical gain Ng.

The luminosity depends on the gain g only through the F_{sto} factor. We can chose an expression of the luminosity which shows the explicit dependence on g as follows:

$$L = f(Ng)L_{opt}, \tag{39}$$

where L_{opt} is the luminosity obtained with the parameters shown in Table 1, that is for the optimum gain $g = 1/N$ and

$$f(u)[(2-u)/u]\exp[-k/(2u-u^2)], \tag{40}$$

with

$$k = 2/(\gamma\tau_0\lambda_{opt}) \ll 1. \tag{41}$$

The first factor of Eq. (41) represents the reduction of the betatron emittance and the exponential factor represents the beam loss. It can be seen that the faction $f(Ng)$ has a maximum for $Ng_{max} \sim k/2$ where it takes the value $f_{max} \sim 1/k$. Thus at most the luminosity can be increased by $f_{max} \sim 0.04\gamma$, that is an increase which is proportional with the beam energy as required. The values of Ng_{max} and of f_{max} with the corresponding increased luminosity are also shown at the bottom part of Table 1. The luminosity figures are still well below the desired values.

The only other parameter that can be varied is the bunching frequency f_{bunc}. For the largest realistic bandwidth, this is equivalent to vary the frequency integer

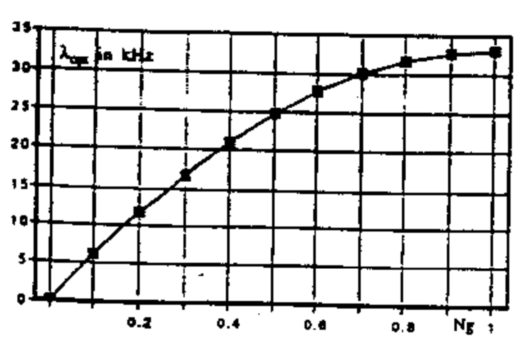

Fig. 6. Cooling rate vs. dynamical gain Ng.

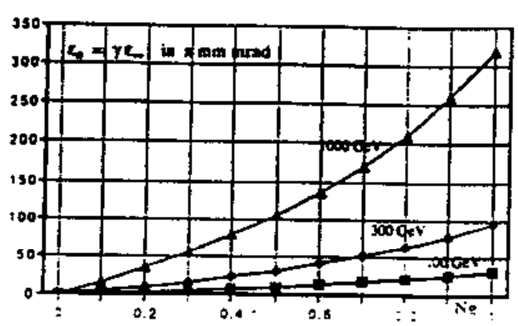

Fig. 8. Equilibrium emittance vs. dynamical gain Ng.

parameter m. If we denote with m_{opt} the value of m used previously, corresponding to $f_{bunc} = 3$ GHz, we can introduce the ratio $\mu = m/m_{opt}$. The expression for the luminosity can then be modified as follows:

$$L = f(Ng, \mu) L_{opt}, \tag{42}$$

where

$$f(u, \mu) = \mu[(2-u)/u] \exp[-\mu k/(2u - u^2)]. \tag{43}$$

The maximum of this function is the one found before and does not depend on the ratio u. That is, the optimum of the luminosity does not depend on the choice of the bunching mode number m, once optimization with respect to the dynamical gain g has been carried out.

9. Conclusions

We have determined that it is indeed feasible that the luminosity of a muon collider scales linearly with the beam energy, as it is required by physics argumentations. Unfortunately, even with the stretching of our imagination, it is seen from our results shown in Table 1 that at the very most only a luminosity of $10^{20}\gamma$ cm^{-2} s^{-1} can be obtained. This is seven orders of magnitude below what it is actually required.

The most important limitation is the effect of thermal noise to the ultimate emittance that can be achieved. This is to be coupled with the requirement on the cooling rate which is to be large compared to the inverse of the beam lifetime. To achieve very fast cooling, a large linear electronic gain is needed, which has also the effect to amplify to a larger level the front-end noise. Moreover a large cooling rate can be obtained only with a few number of particles per bunch. Even by postulating the feasibility of momentum stacking, it is rather difficult to accumulate more than 10^5 particles per bunch.

The comparison of the performance of a muon collider with respect to a proton–antiproton collider is in order. The methodology is essentially the same: both types of particle are produced from a target both need cooling to reduce their dimensions and both are to be compressed longitudinally in bunches. But the antiproton particles have an infinitely long lifetime and it is thus possible to accumulate 10^{10-11} particles per bunch after a long session of stochastic cooling. Moreover it is more convenient for stable particles to operate the collider in a storage mode with multiple passes of the same beams at the collision point.

Our estimates of the performance of stochastic cooling are based on the simple scenario of production, acceleration, cooling and collision we have proposed here. Other scenarios may be possible and we believe that an optimum configuration has still to be searched and is highly desirable. But we also believe that stochastic cooling has to be an integral part of the scheme. Then the following features are to be investigated in more details.

We have seen first that thermal noise at the front-end of the amplifier plays a crucial limiting role to the final beam emittance. We have also seen that this can be better measured with the ratio of Schottky to thermal power, given by Eq. (20). The problem is that as cooling proceeds the beam power reduces whereas the noise signal remains constant. An invention would be highly desirable where the level of noise signal can also be reduced accordingly, as for instance done in momentum cooling with notch filters [9].

The other issue relates to the complete absence of mixing of the particle motion. It is also important to demonstrate that there are ways to regenerate the beam Schottky signal. Several methods have in the meantime been proposed, like the introduction of sextupole to generate a coupling of the transverse motion with the particle momentum error, skew quadrupoles to introduce mixing between the two transverse plane of oscillations, and nonlinear magnetic lenses which cause a dependence of the betatron motion with the amplitude of the motion itself.

These two issues are of paramount importance and are to be investigated carefully if we want to keep the option of a muon collider for the search of the Higgs boson still of some interest.

References

[1] P. Chen and K.T. McDonald, Proc. Workshop on Advanced Accelerator Concepts, AIP Conf. Proc. 279 (1992) 853.

[2] Proc. Mini-Workshop on $\mu^+\mu^-$ colliders: Particle Physics and Design, sponsored by UCLA Center for Advanced Accelerators, Napa Valley, California, Dec. 9–11, 1992.

[3] Proc. Muon Collider Workshop, LA-UR-93-866, LAMPF Auditorium, Feb. 22, 1993.

[4] A.N. Skrinsky, Proc. 20th Int. Conf. on High Energy Physics, AIP Conf. Proc. 68 (1980) 1056.

[5] D. Neuffer, Particle Accel. 14 (1983) 75.

[6] R. Noble, Proc. Workshop on Advanced Accelerator Concepts, AIP, Conf. Proc. 279 (1992) 949.

[7] W.A. Barletta and A.M. Sessler, Characteristics of a high energy $\mu^+\mu^-$ collider based on electro-production of muons, LBL-33613 (1993).

[8] A.G. Ruggiero, Proc. Workshop on Advanced Accelerator Concepts, AIP Conf. Proc. 279 (1992) 958.

[9] S. van der Meer, Physics of Particle Accelerators, AIP Conf. Proc. 153, Volume 2 (1987) 1628.

[10] B. Autin, J. Marriner, A. Ruggiero and K. Takayama, IEEE Trans. Nucl. Sci. NS-30 (4) (1983) 2593.

Nuclear Instruments and Methods in Physics Research A 350 (1994) 53–56
North-Holland

NUCLEAR INSTRUMENTS & METHODS IN PHYSICS RESEARCH
Section A

Critical issues in low energy muon colliders – a summary

S. Chattopadhyay [a], W. Barletta [a], S. Maury [b], D. Neuffer [c], A. Ruggiero [d], A. Sessler [a,*]

[a] Lawrence Berkeley Laboratory, University of California, Berkeley, CA 94720, USA
[b] CERN, CH-1211 Geneva 23, Switzerland
[c] CEBAF, 12000 Jefferson Ave., Newport News, VA 23606, USA
[d] BNL, Upton, New York 11973, USA

Received 13 April 1994

We present a brief summary of the current state of conception and understanding of the accelerator physics issues for low energy muon colliders envisioned as Higgs factories, associated technological challenges and future research directions on this topic.

1. Motivation and challenges

It is well known that multi-TeV e^+–e^- colliders are constrained in energy, luminosity and resolution, being limited by "radiative effects" which scale inversely as the fourth power of the lepton mass ($(E/m_e)^4$). Thus collisions using heavier leptons such as muons offer a potentially easier extension to higher energies [1]. It is also believed that the muons have a much greater direct coupling into the mass-generating "Higgs-sector", which is the acknowledged next frontier to be explored in particle physics. This leads us to the consideration of TeV-scale μ^+–μ^- colliders. However, with the experimental determination of the top quark being heavier than the Z boson, there is increasing possibility of the existence of a "light" Higgs particle with a mass value bracketed by the Z-boson mass and twice that value. This makes a 100 GeV $\mu^+ \otimes$ 100 GeV μ^- collider as a "Higgs Factory" an attractive option [2]. The required average luminosity is determined to be 10^{30} cm^{-2} s^{-1} [2]. We note that the required luminosity for the same "physics reach" scales inversely as the square of the lepton mass and implies a significantly higher luminosity required of a similar energy e^+–e^- collider, in order to reach the same physics goals.

The challenges associated with developing a muon collider were discussed at the Port Jefferson workshop [1,3], subsequent mini-workshops at Napa [2], Los Alamos [4] and at the workshop [5,6] on "Beam Cooling and Related Topics", in Montreux, Switzerland in 1993. Basically, the two inter-related fundamental aspects about muons that critically determine and limit the design and development of a muon collider are that muons are secondary particles and that they have a rather short lifetime in the rest frame. The muon lifetime is about 2.2 μs at rest and is dilated to about 2.2 ms at 100 GeV in the laboratory frame by the relativistic effect. The dilated lifetime is short enough to pose significant challenges to fast beam manipulation and control. Being secondary particles with short lifetime, muons are not to be found in abundance in nature, but rather have to be created in collisions with heavy nuclear targets. Muon beams produced from such heavy targets have spot size and divergence-limited intrinsic phase-space density which is rather low. To achieve the require luminosity, one needs to cool the beams in phase-space by several orders of magnitude. And all these processes – production, cooling, other bunch manipulations, acceleration and eventual transport to collision point – will have to be completed quickly, in 1–2 ms, and therein lies the challenge. Bunch manipulation and cooling of phase space are some of the primary concerns. In the following section, we describe the two scenarios, and associated parameters being considered at present for muon colliders.

2. Scenarios, parameters and comments

Basically, there are two scenarios that have been considered to date for muon colliders. These two scenarios start with very different approaches to the production of the secondary muon beam from a primary beam hitting a heavy target. The subsequent acceleration, cooling, stacking, bunching and colliding gymnastics are all dictated and differentiated by these production schemes, which are very different. We consider them in sequence in the following.

The first approach considers production of the muons starting from a primary "proton" beam hitting a heavy target according to the following reaction:

$$p + N \rightarrow \pi + X, \quad \pi \rightarrow \mu\nu.$$

* Corresponding author.

0168-9002/94/$07.00
SSDI 0168-9002(94)00540-N

Since proton bunches are typically long (few ns), one basically obtains long bunches of low phase-space density unless further phase-space manipulations are done to bunch and cool the beams. The situation is similar to the use of the Proton Ring as a pion source in the Los Alamos Meson Physics Facility (LAMPF-II) or conventionally considered kaon factory sources, for example. In order to reduce the length of the produced muon beam bunches, considerable gymnastics is required of the proton ring rf system. Ultimately, of course, a bunch rotation in the longitudinal phase space to reduce bunch length comes at the expense of the relative momentum spread, $\Delta p/p$, which could be as high as 5%. The produced muon bunches will need to be cooled longitudinally from $\Delta p/p$ of 5% to about 0.1% in order to have acceptable spectral purity at the collision point. In addition, the muon bunches will have to be cooled in the transverse phase space by a significant amount in order to meet the luminosity demand at the collision point. The cooled muons are subsequently accelerated and injected into a 100 GeV μ^+–μ^- collider where the bunches collide in at most a few hundred to a thousand turns (the number of turns, $n \approx 300\bar{B}$ [T]). Clearly the constraint of short muon lifetime puts a premium at every stage on minimizing the time for production, cooling, acceleration and bunch processing, so as to still leave a few hundred turns in the collider to produce luminosity. Thus, it is clear that high field magnets play a crucial role in the collider. Details of this scenario have been considered by Neuffer [2,5]. In Fig. 1, we depict schematically the scenario of a muon collider based on production via protons [5].

A second approach considers production of the muons starting from a primary "electron" beam hitting a heavy target according to the following reaction:

$$e + N \rightarrow e + N + \gamma$$
$$\qquad\qquad \hookrightarrow \mu^+ \mu^- .$$

In this electro-production scenario, one obtains short bunches most naturally, since it is compatible with the normal mode of operation of high energy linacs. Although one obtains the "optimum bunch format" naturally, one has to consider unprecedently high power and high repetition rate electron linacs, not explored before in order to meet the required collision luminosity. This is so because of the rather low yield of muons per electron, even at the optimum energy of incident electrons of 60 GeV, and the difficulty of packing more electrons per bunch in the linac. The low transverse phase-space density of the muons will require significant improvement via cooling, similar to the proton production scenario, and, in addition, calls for a nontrivial beam stacking scheme before collision (described in Ref. [6]). Details of this scenario have been considered by Barletta and Sessler [2,6]. In Fig. 2, we depict schematically the scenario of a muon collider based on electro-production [6].

Table 1 presents a comparison of parameters for the above two scenarios for a 100 GeV $\mu^+ \otimes$ 100 GeV μ^- collider, with an average luminosity of 10^{30} cm^{-2} s^{-1}. We assume a collider scenario with a low beta at the collision point of 1 cm, about 1000 bunches colliding in

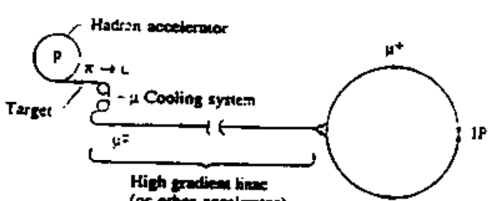

Fig. 1. Overview of a μ^+–μ^- collider, showing a hadronic accelerator, which produces π's on a target, followed by a μ-decay channel ($\pi \rightarrow \mu\nu$) and μ-cooling system, followed by a μ-accelerating linac (or recirculating linac or rapid-cycling synchrotron), feeding into a high-energy storage ring for μ^+–μ^- (from Ref. [5]).

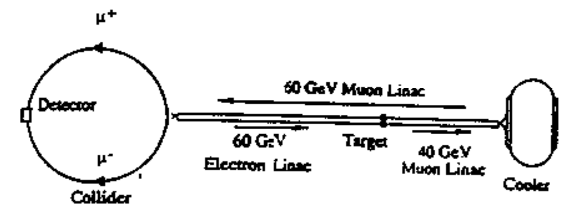

Fig. 2. A muon collider scenario with electro-production of muons (from Ref. [6]).

Table 1
Parameters for a muon collider, 100 GeV × 100 GeV,
$$L = M\frac{N_+ N_- f}{4\pi\epsilon_N \beta^*}\gamma \sim 10^{30}\ \text{cm}^{-2}\ \text{s}^{-1},$$
$M = 1000$; $\gamma = 1000$; $\beta^* = 1$ cm; $P = 5$ MW at the target

	Production via electrons	Production via protons
$E_{e\ \text{or}\ p}$ (GeV)	60	30
Intensity	5×10^{11} /pulse	10^{14}/pulse
Number of pulses	100 (stacked later)	1
Repetition rate	10 Hz	10 Hz
E_μ (GeV)	40	1.5
ϵ_N (π m-rad)	2×10^{-3}	2×10^{-2}
$\Delta p/p$	±3%	±3%
(μ/e) or (μ/p)	4×10^{-3}	10^{-3}
Ionization cooling	$\epsilon_n^f = 2\times10^{-5}\pi$ m-rad	$\epsilon_n^f = 2\times10^{-5}\pi$ m-rad
Bunch rotation factor	none	100

the ring and muon production limited by a 5 MW power at the target. It is clear that while powerful pion sources, bunch compression and cooling are essential for the proton-production scenario, high current electron linacs, cooling and stacking are essential for the electro-production scenario. It is fair to say from an inspection of Table 1 that, fundamentally, both scenarios are equally amenable to a muon collider configuration with comparable luminosities, given the fact that in both cases equally difficult and challenging technological problems will have to be addressed and solved.

The most difficult and challenging of these technological problems is probably that of "ultra-rapid" phase space cooling of "intense" bunches. One can consider radiation cooling via synchrotron radiation, which is independent of the bunch intensity. However, it is too slow for our purposes. The stochastic cooling rate, on the other hand, depends on the number of particles per bunch and, although too slow usually, can be made significantly faster by going to an extreme scenario of a few particles per bunch with ultra-fast phase mixing or an ultra-high bandwidth ($\sim 10^{14}$ Hz) cooling feedback loop. Both the latter cases will require significant technological inventions. A promising scheme that is both "fast" and "intensity-independent" is that of "ionization cooling", which looks feasible in principle. We have assumed ionization cooling in arriving at the parameters of Table 1. We discuss cooling considerations briefly in the next section.

3. Cooling of muons

The cooling of the transverse phase-space assumed in Table 1 is of the kind known as "ionization cooling". In this scheme the beam transverse and longitudinal energy losses in passing through a material medium are followed by coherent reacceleration, resulting in beam phase-space cooling [2,5,7]. The cooling rate achievable is much faster than, although similar conceptually to, radiation damping in a storage ring in which energy losses in synchrotron radiation followed by rf acceleration result in beam phase-space cooling in all dimensions. Ionization cooling is described in great detail in Refs. [2,5]. It seems that the time is ripe to make a serious design of an ionization cooling channel, including the associated magnetic optics and rf aspects, and put it to real test at some laboratory.

Exploration of the alternate cooling scheme of stochastic cooling takes us to a totally different regime of operation of the collider, determined by the very different nature and mechanism of cooling by an electronic feedback system. Here, the muon lifetime and the required low emittance demanded by the luminosity requirements determine the necessary stochastic cooling rate of the phase space. This rate scales directly as the bandwidth (W) of the feedback system and inversely as the number of particles (N) in the beam (stochastic cooling rate $\propto W/N$). If we limit our consideration to practically achievable conventional feedback electronics, amplifiers, etc., with bandwidth not exceeding 10 GHz, the number of particles per bunch must be less than a thousand in order to meet the desired rate. This then would imply a very different pulse format. This alone drives all the parameters back to the source and issues of "targetry" and "muon source", etc., are not critical. The critical issues for stochastic cooling are: (1) large bandwidth, (2) ultra-low noise, as the cooled emittance reaches the thermal limit of the electronics, (3) rapid mixing and (4) bunch recombination techniques.

Critical issues in the stochastic cooling scenario are discussed by Ruggiero [2,8], where he also explores a conventional cooling scheme with modest bandwidth but with a special nonlinear (magnetic) device that stirs up the phase space rapidly and provides "ultra-fast mixing". It is clear that we need new technical inventions in stochastic cooling for application in a muon collider. Another novel scheme [9] being explored currently is that of "optical cooling" where one detects the granularity of phase space down to a micron scale by carefully monitoring the incoherent radiation from the beam, which is a measure of its Schottky noise, then amplifying this radiation via a laser amplifier of high gain and bandwidth (10^7, 100 THz) and applying it back to the beam. Various issues regarding quantum noise and effective pickup and kicker mechanisms will have to be understood before it can be considered for a serious design.

4. Summary and outlook

As we have seen, both scenarios – production of muons from protons and electro-production of muons – are competitive but very ambitious and challenging. Production of muons from protons will clearly require nontrivial and sophisticated target design and configuration. In addition, in order to match the bunch length of the colliding (but secondarily produced) muon beams to the low beta function at the collision point, the primary proton beams must be bunched by a large factor (~ 100). The complicated bunch rotation and rf manipulations are cumbersome and must be done at the low energy proton end before the target, which implies an associated increase in the relative momentum spread, $\Delta p/p$. On a positive note, however, targetry with protons and rf gymnastics with proton beams are relatively familiar affairs at hadron and kaon facilities, albeit at a lower level of power and rf manipulation of the bunches. Electro-production of muons, on the other hand, requires, high peak current, high repetition rate linacs, so far unexplored, in order to meet the luminosity demand. Besides, "stacking" of many electron bunches from a linac into a single bunch poses a nontrivial problem. The significant and most attractive feature of the electro-production scenario, however, is that the "optimal pulse format" is produced directly at the target by electrons from a linac, without complex bunch compression schemes in a ring.

No matter what the optimal scenario would turn out to be, should the muon collider concept turn into reality, further consideration of such a collider at 200 GeV center-of-mass energy with an average luminosity of $\sim 10^{30}$ cm^{-2} s^{-1} would have to assume major advances in, and eventual operation of, (1) megawatt muon targets, (2) multi-kiloampere peak current electron linacs, (3) efficient transfer, compression and stacking schemes for charged particle beams, (4) high field magnets and (5), most importantly, feasible phase-space cooling technologies with low noise and large bandwidth. While "ionization cooling" looks promising, it needs experimental demonstration. A possible feasibility test of muon production and ionization cooling at existing facilities, e.g., CERN or FNAL, would be highly desirable. The "stochastic cooling" approach, however, would need fundamental invention of a new technique, as elaborated earlier. The emerging new ideas of "optical stochastic cooling", "ultra-rapid phase-mixer", etc., are ambitious, but may hold the key to the success of such high frequency stochastic cooling. Finally, the synchrotron radiation and muon decay in the collider ring vacuum chamber and detector area pose issues that cannot be overlooked.

In conclusion, surely a muon collider is exotic! But even as we contemplate the value, utility and eventual realizability of such a collider in the future, there is no doubt that the necessary conceptual and technological explorations forced upon us by these considerations are much too valuable to many fields to be simply passed up.

Acknowledgement

This work was supported by the Director, Office of Energy Research, Office of High Energy and Nuclear Physics, High Energy Physics Division, of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

References

[1] M. Tigner, Proc. Advanced Accelerator Concepts, Port Jefferson, NY, 1992, AIP Conf. Proc. 279 (1993) 1.
[2] D.B. Cline, this issue, Nucl. Instr. and Meth. A 350 (1994) 24.
[3] P. Chen and K.T. MacDonald, AIP Conf. Proc. 279 (1993) 853.
[4] Proc. Muon Collider Workshop, February 22, 1993, Los Alamos National Laboratory Report LA-UR-93–866 (1993).
[5] D. Neuffer, Muon Cooling and Applications, Proc. of the Workshop on Beam Cooling and Related Topics, October 3–8, 1993, Montreux, Switzerland, to be published as a CERN report.
[6] W.A. Barletta and A.M. Sessler, Prospects for a High Energy Muon Collider Based On Electro-production, Proc. Workshop on Beam Cooling and Related Topics, October 3–9, 1993, Montreux, Switzerland, to be published as a CERN report.
[7] D. Neuffer, Particle Accel. 14 (1983) 75.
[8] A.G. Ruggiero, Stochastic Cooling Requirements for a Muon Collider, Proc. Workshop on Beam Cooling and Related Topics, October 3–9, 1993, Montreux, Switzerland, to be published as a CERN report.
[9] A.A. Mikhailichenko and M.S. Zolotorev, Optical Stochastic Cooling, SLAC-PUB- 6272 (1993).